Efficient magnetic fields for supporting toroidal plasmas

Matt Landreman*
Institute for Research in Electronics and Applied Physics, University of Maryland, College Park, MD, 20742, USA

Allen Boozer
Department of Applied Physics and Applied Mathematics, Columbia University, New York, NY, 10027, USA
(Dated: March 8, 2016)

The magnetic field that supports tokamak and stellarator plasmas must be produced by coils well separated from the plasma. However, the larger the separation, the more difficult it is to produce a given magnetic field in the plasma region, so plasma configurations should be chosen that can be supported as efficiently as possible by distant coils. The efficiency of an externally-generated magnetic field is a measure of the field's shaping component magnitude at the plasma compared to the magnitude near the coils; the efficiency of a plasma equilibrium can be measured using the efficiency of the required external shaping field. Counterintuitively, plasma shapes with low curvature and spectral width may have low efficiency, whereas plasma shapes with sharp edges may have high efficiency. Two precise measures of magnetic field efficiency, which correctly identify such differences in difficulty, will be examined. These measures, which can be expressed as matrices, relate the externally-produced normal magnetic field on the plasma surface to the either the normal field or current on a distant control surface. A singular value decomposition (SVD) of either matrix yields an efficiency ordered basis for the magnetic field distributions. Calculations are carried out for both tokamak and stellarator cases. For axisymmetric surfaces with circular cross-section, the SVD is calculated analytically, and the range of poloidal and toroidal mode numbers that can be controlled to a given desired level is determined. If formulated properly, these efficiency measures are independent of the coordinates used to parameterize the surfaces.

arXiv:1601.05741v2 [physics.plasm-ph] 6 Mar 2016

* mattland@umd.edu

I. INTRODUCTION

In schemes for magnetic plasma confinement, including both tokamaks and stellarators, one must create and control magnetic fields in the plasma region using electromagnetic coils at a distance from the plasma. In tokamaks, not only is it typically beneficial to provide axisymmetric shaping (elongation, triangularity, an X-point, etc), but also to manipulate nonaxisymmetric magnetic field components in order to control error fields and edge localized modes (ELMs). In stellarators, coils shapes are optimized to achieve a plasma shape that has good physics properties. In both tokamaks and stellarators, the distance between the coils and plasma is a critical quantity. Nearby coils cannot be used in a nuclear environment, and a small coil-plasma separation can complicate design and construction, increasing costs. (Indeed, much of the challenge in assembling the W7-X stellarator stemmed from the need to fit many device components in the small space between plasma and coils, "lesson 1" in [1].) The distance between coils and plasma can sometimes be increased at the cost of greater current in the coils and/or greater coil curvature, but this increase leads to other engineering challenges such as greater electromagnetic forces in the coils. However, beyond a certain plasma-coil separation, it is typically impossible to find practical coils to support a given plasma shape within a reasonable tolerance [2]. Thus, magnetic confinement fusion program would greatly benefit from improved and systematic understanding of the efficient control of magnetic fields using distant coils.

Efficiency is a property of both external magnetic fields (those generated by currents outside the plasma) and of plasma equilibria. For an external magnetic field, efficiency is a measure of the field's shaping component magnitude at the plasma divided by some measure of the magnitude far from the plasma. The efficiency of a plasma equilibrium can then be measured by the efficiency of the external field required to support the plasma (in the sense of a free-boundary equilibrium). These rough definitions will be made precise below.

One might naively expect that the efficiency of a given plasma equilibrium is directly related to the minimum spatial length scale of the plasma boundary shape, with structures of length scale λ being inefficient to produce using currents displaced from the plasma by a length $\geq \lambda$. However, this reasoning is too simplistic, as the following counterexample shows. A diverted tokamak has a corner at the X-point, meaning the plasma shape contains arbitrarily small length scales. Nonetheless, it has been feasible to produce the X-point shape in many tokamak facilities, and in the design of ITER and many reactor concepts. The divertor coil current required to produce an X-point, which for ITER is $\leq 1.5\times$ the plasma current, has not been prohibitive in practice, so the infinitesmal spatial scales λ evidently do not imply an infinitesmal efficiency. Another quantity which fails to measure the efficiency of a plasma equilibrium is the spectral width [3] of the boundary shape, which we define loosely as the number of Fourier modes needed to represent the shape via $R(\theta,\zeta) = \sum_{m,n} R_{m,n} e^{i(m\theta - n\zeta)}$ and $Z(\theta,\zeta) = \sum_{m,n} Z_{m,n} e^{i(m\theta - n\zeta)}$ for some poloidal and toroidal angles θ, ζ. Many Fourier modes are required to represent the boundaries of tokamak plasmas with divertors, while other shapes with very small spectral width, such as shapes with concavities, can be difficult to create with external coils. Furthermore, the spectral width depends on the specific choice of coordinates. Better measures of efficiency are needed.

Stellarator optimizations to date have typically been based upon using a small number of Fourier modes to describe the plasma shape [4], This space of shape parameters may exclude some desirable plasma configurations with large spectral width and/or sharp corners, although this is not certain - one can drop the outer surfaces of such an equilibrium to find a smaller similar equilibrium within the parameter space. This space of shape parameters will also include regions where it is difficult to find reasonable coils. To address this issue, recent optimizations have included measures of coil complexity in the overall target function[5]. In such a procedure, measures should be chosen that do not necessarily penalize sharp edges, since such shapes are not necessarily inefficient to produce, and sharp edges may indeed be advantageous for a divertor.

Here, we will examine and compare two superior possible definitions of magnetic field efficiency, the "transfer operator" efficiency

$$\text{Efficiency} = \frac{B_n \text{ on plasma surface}}{B_n \text{ on a distant control surface}} \tag{1}$$

and the "inductance operator" efficiency

$$\text{Efficiency} = \frac{B_n \text{ on plasma surface}}{\text{Current potential on a distant control surface}}, \tag{2}$$

where B_n denotes the magnetic field component normal to the surface, and the current potential is the stream function for a divergence-free skin current. The control surface in these definitions is a toroidal surface offset by some distance (uniform or nonuniform) from the plasma. Depending on the application, the control surface may be either a specific surface on which we wish to design coils, or any surface that is relatively far from the plasma.

Our definitions emphasize B_n because this is the 'shaping component', the relevant component for controlling the plasma surface shape. For any fixed-boundary MHD equilibrium solution, the same equilibrium could be realized in

the realistic scenario of a free boundary if coils could be designed to achieve $B_n = 0$ on the desired plasma boundary [2]. The linearity of Ampère's Law means we can write $B_n = B_n^{pl} + B_n^{fix} + B_n^x$ as a sum of contributions B_n^{pl} from current in the plasma, B_n^{fix} from fixed coil currents outside the plasma, and B_n^x from other external shaping currents which are being determined. The B_n^{fix} term may represent the field due to standard planar toroidal field coils, and/or the net poloidal current on a modular coil winding surface which creates the main toroidal field. The B_n^x term may represent the field due to saddle coils, nonplanar excursions of modular coils, poloidal shaping coils, etc. (Our division of the externally generated B_n into B_n^{fix} and B_n^x corresponds to Merkel's division into h and g terms in [2].) The problem of designing shaping coils is thus to find shaping currents that achieve $B_n^x \approx -B_n^{pl} - B_n^{fix}$, hence our focus on controlling B_n.

The two notions of efficiency (1)-(2) arise from two linear operators, the transfer and inductance operators. The inductance operator relates skin currents on the control surface to B_n on the plasma surface. This approach is closely related to the `NESCOIL` method of stellarator coil design [2]. The transfer operator instead relates B_n on the control surface to B_n on the plasma surface. For the part of the magnetic field driven by external coils, $\boldsymbol{B} = \nabla\phi$ for some scalar potential ϕ satisfying $\nabla^2\phi = 0$. A given $B_n = \boldsymbol{n} \cdot \nabla\phi$ on the control surface (with normal vector $\boldsymbol{n}$) represents a Neumann boundary condition, so Laplace's equation $\nabla^2\phi = 0$ inside the control surface has a unique solution up to a constant, and so a unique B_n on the plasma surface can be determined. The transfer operator encodes this relationship between B_n on the two surfaces. Finite-dimensional inductance and transfer matrices arise from discretizing the quantities of interest on the two surfaces by expanding in a finite number of basis functions.

For both the inductance and transfer matrices, a quantitative measure of magnetic field efficiency can be obtained using a singular value decomposition (SVD). The SVD results in an orthonormal set of right singular vectors, each representing a pattern of B_n or the current potential on the control surface, paired with a left singular vector representing the resulting pattern of B_n on the plasma surface. Each pair of left and right singular vectors is associated with a singular value, which measures the efficiency by which currents or B_n on the control surface drive B_n on the plasma surface. Each singular value, with the associated right and left singular vectors, represents a vacuum solution of $\nabla^2\phi$ throughout the volume enclosed by the control surface, and hence a pattern of $\boldsymbol{B}(\boldsymbol{r})$ everywhere in this volume, not just on the two surfaces. We call each of these spatial patterns of $\boldsymbol{B}$ (as well as the associated B_n on the surfaces) a 'magnetic field distribution'.

Given the transfer or inductance matrix SVD and $B_n^{pl} + B_n^{fix}$, the efficiency of a plasma equilibrium shape can be measured using either of two sequences, which we call the efficiency and feasibility sequences. The efficiency sequence is the projection of $B_n^{pl} + B_n^{fix}$ onto the left singular vectors. It has an exponentially decreasing trend, with rapid decrease for plasma shapes that can be efficiently produced. For example, we will show this decrease is rapid for a diverted tokamak shape, despite the sharp corner at the X point. This sequence is insensitive to the particular choice of control surface, so its rate of decrease measures the intrinsic efficiency of the plasma shape. The feasibility sequence is the efficiency sequence divided by the singular values, representing the amplitudes of B_n or current on the control surface required to achieve $B_n = 0$ on the plasma surface. If this sequence is (exponentially) decreasing, the required amplitudes on the control surface are bounded, so it is feasible to support this plasma shape with coils on this control surface. A plasma shape could be optimized by targeting the rate of decrease of either the efficiency or feasibility sequence.

We examine both the inductance and transfer matrices since each has advantages and disadvantages. The transfer matrix relates a B_n to a B_n, so both the matrix and its singular values are naturally dimensionless. Its singular values have the satisfying property that they approach 1 as the plasma and control surfaces approach each other. In contrast, the inductance matrix is not dimensionless, and its singular values do not have a simple limiting value when the surfaces coincide. However, ultimately we are interested in generating the magnetic field with currents on the control surface, and the connection to the generating current is clearer for the inductance matrix than for the transfer matrix. The inductance matrix is easier and faster to compute, since computation of the transfer matrix requires inverting a dense linear system, and one must carefully integrate over certain singularities (as is common in boundary integral equation methods.) We will find that the inductance and transfer matrices generally give a similar description of external control of the magnetic field. Some long-wavelength structures are found to be more efficient using the transfer matrix than using the inductance matrix. The reason is that magnetic field on the plasma surface is generated by gradients of the current potential rather than by the current potential itself; long-wavelength magnetic field structures can be inefficiently generated even though they propagate efficiently.

Other measures of magnetic field efficiency could be considered besides the ones we consider here. One example is the radius of curvature of the conductors generating the shaping fields, related to curvature of contours of the current potential. Another example is the magnetic energy in the region enclosed by the plasma bounding surface divided by the magnetic energy throughout space.

The present work builds upon the work of several previous authors. SVD techniques have been applied to an inductance matrix in the extension of `NESCOIL` by Pomphrey et al [5], and are discussed in [6, 7]. The concept of the transfer matrix was introduced by Pomphrey in unpublished work, and his numerical calculation appears in figure

2 of Ref [8]. The concept of the transfer matrix and the importance of its singular value decomposition were also discussed elsewhere in Ref [8] and in [9].

In the following two sections we introduce preliminary concepts, then review the definitions of the inductance matrix and transfer matrix. Some details of our numerical implementation are given in section IV. In section V we build some intuition for the two matrices and their SVDs by considering the case of axisymmetric surfaces with circular cross-section, since in this limit the SVDs can be computed analytically using a large-aspect-ratio approximation. We also demonstrate that these analytical results agree extremely well with finite-aspect-ratio numerical calculations. Shaped axisymmetric surfaces are examined in section VI, and here we demonstrate (figure 10) that both the inductance and transfer matrix find the X-point divertor shape to be efficient, despite the sharp corner. A numerical example for the W7-X stellarator is presented in section VII. Here we show that our efficiency measures are consistent with the choice of coil winding surface for W7-X, that it is as far from the plasma as possible while still allowing the desired plasma shape to be supported. In section VIII, we prove that for a suitable definition of orthogonality and within discretization error, the SVD is independent of the coordinates chosen to parameterize the surfaces, and also independent of the specific choice of basis functions. We conclude in section IX. An alternative method for visualizing the relationship between the inductance and transfer matrices is presented in the appendix.

II. CURRENT POTENTIAL AND BASIS FUNCTIONS

Divergence-free currents on a surface have the form

$$\boldsymbol{K} = \boldsymbol{n} \times \nabla \kappa \tag{3}$$

where $\boldsymbol{K}$ denotes a surface current density, $\boldsymbol{n}$ is the outward normal vector associated with the surface, and the quantity κ is known as the current potential [2]. Contours of the current potential represent streamlines of the current.

In general the current potential may be multiple-valued, because there may be net current that topologically links the surface. For example, the standard toroidal field coils in a tokamak correspond to a current potential that increases as one moves toroidally around the surface, $\kappa \propto \zeta$ where ζ is the toroidal coordinate. In general we can write

$$\kappa = \kappa_{sv} + \frac{G\zeta}{2\pi} + \frac{I\theta}{2\pi}, \tag{4}$$

where G and I are the total currents linking the surface poloidally and toroidally, θ and ζ are any poloidal and toroidal angle coordinates, and κ_{sv} is single-valued. Since G and I are fixed given the enclosed toroidal flux and the choice of current topology (corresponding to saddle, modular, or helical coils) [2, 14], we can consider the magnetic field driven by $G\zeta/(2\pi) + I\theta/(2\pi)$ to be part of the 'known' B_n^{fix}, whereas κ_{sv} generates B_n^x and is considered 'unknown', adopting the same philosophy as in `NESCOIL`. Hence, for the rest of this paper we will be concerned only with the single-valued part κ_{sv}, dropping the subscript to simplify notation.

To describe scalar quantities on surfaces, such as the current potential and the normal component of $\boldsymbol{B}$, we will expand in some set of basis functions $f_j(\theta, \zeta)$. Each surface generally may have its own set of basis functions. It is convenient to require that the basis functions satisfy an orthogonality relation

$$\int \mathrm{d}^2 a \; w f_i f_j = \delta_{i,j} \tag{5}$$

where $\int \mathrm{d}^2 a$ represents an area integral over the surface, and $w(\theta, \zeta)$ is some weight, normalized so that

$$\int \mathrm{d}^2 a \; w = 1. \tag{6}$$

As we will see below, it can sometimes be useful to consider non-constant w. We write the normal component of the magnetic field B_n as a weighted linear combination of the basis functions:

$$B_n = w \; \Phi_j f_j \tag{7}$$

for some components Φ_j. (Repeated indices imply summation.) Using (5),

$$\Phi_j = \int \mathrm{d}^2 a \; B_n f_j = \int \mathrm{d}^2 \boldsymbol{a} \cdot \boldsymbol{B} \; f_j. \tag{8}$$

The components Φ_j may be thought of as magnetic fluxes, as they are area integrals of the magnetic field. Similarly, the single-valued part of the current potential may be written

$$\kappa = I_j f_j \tag{9}$$

for some components I_j, satisfying

$$I_j = \int \mathrm{d}^2 a \; w \kappa f_j. \tag{10}$$

Both κ and I_j have units of Ampères. Notice that there is some freedom of definition in whether to include the weight w in (7) or (8), and in (9) or (10). The choices used in (7)-(10) are made for consistency with [8], where they enable several analogies to be drawn between toroidal magnetic fields and electric circuits.

One natural choice of basis functions is the Fourier basis, consisting of sines and cosines in the poloidal and toroidal angles. This basis requires a specific choice of weight w. Observe that area integrals of any quantity Q may be written

$$\int \mathrm{d}^2 a \; Q = \sum_{\ell=1}^{n_p} \int_0^{2\pi} \mathrm{d}\theta \int_0^{2\pi/n_p} \mathrm{d}\zeta \; NQ \tag{11}$$

where n_p is the number of identical toroidal periods (e.g. $n_p = 5$ for W7-X.) Also, $N = |\boldsymbol{N}|$ where

$$\boldsymbol{N} = \frac{\partial \boldsymbol{r}}{\partial \zeta} \times \frac{\partial \boldsymbol{r}}{\partial \theta} \tag{12}$$

is the (non-unit-length) normal vector. In order for trigonometric functions to satisfy the orthogonality relation (5), we must choose $w \propto 1/N$ to cancel the factor of N in (11). The choice

$$w = 1/(4\pi^2 N) \tag{13}$$

is then required by (6). Then from (5), the functions

$$\begin{aligned} f_j(\theta, \zeta) = \sqrt{2} \sin(m_j \theta - n_j \zeta) \text{ and} \\ \sqrt{2} \cos(m_j \theta - n_j \zeta), \end{aligned} \tag{14}$$

for integers m_j and n_j are valid basis functions satisfying our requirements. (We exclude the constant basis function with $m = n = 0$. The component Φ_j associated this basis function must have zero amplitude due to $\nabla \cdot \boldsymbol{B} = 0$ and (8). We can also require the amplitude of the constant basis function to be zero for the expansion of the current potential I_j, since we can shift κ by an arbitrary constant without altering the physical surface current.) Stellarator symmetry (the 3D generalization of tokamak up-down symmetry) can be imposed conveniently by considering only the sine functions in (14).

As we will discuss in section VIII, the Fourier basis (14) has an unsatisfying property that some important quantities we will compute depend on the choice of coordinates θ and ζ used to parameterize the surfaces. For example, the plasma surface could be parameterized by the coordinates of the widely-used `VMEC` code [10] (in which magnetic field lines are not straight), or any of the popular straight-field-line coordinates. In section VIII we will show this coordinate-dependence is a result of the fact that the weight (13) depends on the choice of coordinates. Therefore, it is advantageous to consider a different set of basis functions orthogonal under a weight that is constant on the surface, and therefore independent of the choice of coordinates. From (6), such a weight must be $w = 1/A$ where $A = \int \mathrm{d}^2 a$ is the area of the surface. There are multiple ways to generate suitable basis functions. One approach is to use the functions

$$\begin{aligned} f_j(\theta, \zeta) = \sqrt{\frac{A}{2\pi^2 N(\theta, \zeta)}} \sin(m_j \theta - n_j \zeta) \text{ and} \\ \sqrt{\frac{A}{2\pi^2 N(\theta, \zeta)}} \cos(m_j \theta - n_j \zeta) \end{aligned} \tag{15}$$

for integers m_j and n_j. This is the basis set used for the rest of this paper, except where noted otherwise. A different set of basis functions orthogonal under the same constant weight $w = 1/A$ is discussed in section VIII.

III. INDUCTANCE AND TRANSFER MATRICES

Consider a toroidal surface on which a skin current $\boldsymbol{K}$ flows. We can compute the resulting magnetic field at any point using the Biot-Savart Law

$$\boldsymbol{B}(\boldsymbol{r}) = \frac{\mu_0}{4\pi}\int \mathrm{d}^2a' \frac{\boldsymbol{K}(\boldsymbol{r}')\times(\boldsymbol{r}-\boldsymbol{r}')}{|\boldsymbol{r}-\boldsymbol{r}'|^3}. \tag{16}$$

Primes denote quantities on the surface. Substituting in (3), (11), and (12),

$$\boldsymbol{B}(\boldsymbol{r}) = \frac{\mu_0}{4\pi}\sum_{\ell'=1}^{n_p}\int_0^{2\pi}\mathrm{d}\theta'\int_0^{2\pi/n_p}\mathrm{d}\zeta'\left(\frac{\partial\kappa'}{\partial\zeta'}\frac{\partial\boldsymbol{r}'}{\partial\theta'}-\frac{\partial\kappa'}{\partial\theta'}\frac{\partial\boldsymbol{r}'}{\partial\zeta'}\right)\times\frac{(\boldsymbol{r}-\boldsymbol{r}')}{|\boldsymbol{r}-\boldsymbol{r}'|^3}. \tag{17}$$

Integrating by parts in θ' and ζ' to remove the derivatives on κ', after some algebra we obtain

$$\boldsymbol{B}(\boldsymbol{r}) = \frac{\mu_0}{4\pi}\int \mathrm{d}^2a\ \kappa'\left[\frac{1}{|\boldsymbol{r}-\boldsymbol{r}'|^3}\boldsymbol{n}'-\frac{3}{|\boldsymbol{r}-\boldsymbol{r}'|^5}\boldsymbol{n}'\cdot(\boldsymbol{r}-\boldsymbol{r}')(\boldsymbol{r}-\boldsymbol{r}')\right], \tag{18}$$

where $\boldsymbol{n} = \boldsymbol{N}/|\boldsymbol{N}|$. This magnetic field is precisely that generated by magnetic dipoles of density κ covering the surface and oriented normal to it.

Let us call the surface considered so far the "control" surface outside the plasma, and consider a second toroidal surface inside this control surface, representing the boundary of a plasma. We evaluate the component of (18) normal to this inner surface, giving a linear relation between κ and the magnetic field normal to the plasma surface, denoted B_n^p. The operator M_{op} in this linear relation $B_n^p = M_{op}\kappa$ is the inductance operator. A discretized linear relation is obtained by applying the operation $\int \mathrm{d}^2a f_i^p(\theta,\zeta)(\ldots)$, where $f_i^p(\theta,\zeta)$ are the basis functions on the plasma surface. Recalling (8), (9), and (10), and letting $f_j^c(\theta',\zeta')$ denote the basis functions on the control surface, we obtain

$$\Phi_i = M_{i,j}I_j \tag{19}$$

(or in alternate notation $\boldsymbol{\Phi} = \mathbf{M}\boldsymbol{I}$) where

$$M_{i,j} = \frac{\mu_0}{4\pi}\int \mathrm{d}^2a\int \mathrm{d}^2a'\ f_i^p(\theta,\zeta)f_j^c(\theta',\zeta')\left[\frac{\boldsymbol{n}\cdot\boldsymbol{n}'}{|\boldsymbol{r}-\boldsymbol{r}'|^3}-\frac{3(\boldsymbol{r}-\boldsymbol{r}')\cdot\boldsymbol{n}(\boldsymbol{r}-\boldsymbol{r}')\cdot\boldsymbol{n}'}{|\boldsymbol{r}-\boldsymbol{r}'|^5}\right]. \tag{20}$$

Quantities without primes refer to the plasma surface. We call $M_{i,j}$ the (mutual) inductance matrix between the two surfaces, since it relates currents to fluxes.

Next, we define the transfer operator and matrix. As noted in the introduction, given B_n on the control surface, denoted B_n^c, Laplace's equation $\nabla^2\phi = 0$ gives a unique solution for B_n on the plasma surface, denoted B_n^p. Due to the linearity of Laplace's equation, B_n^c is a linear functional of B_n^p: $B_n^p = T_{op}B_n^c$ for some linear operator T_{op}. We call T_{op} the transfer operator. A discrete version of this linear relation also holds: the vector of fluxes on the plasma surface $\boldsymbol{\Phi}^p$ must be a linear function of the vector of fluxes on the control surface $\boldsymbol{\Phi}^c$. The transfer matrix $\mathbf{T}$ is precisely this linear relationship:

$$\boldsymbol{\Phi}^p = \mathbf{T}\boldsymbol{\Phi}^c. \tag{21}$$

For any real matrix $\mathbf{A}$, including the inductance and transfer matrix, one can find a singular value decomposition

$$\mathbf{A} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T \tag{22}$$

where $\mathbf{U}$ and $\mathbf{V}$ are orthogonal ($\mathbf{U}^{-1} = \mathbf{U}^T$ and $\mathbf{V}^{-1} = \mathbf{V}^T$), $\boldsymbol{\Sigma}$ is diagonal, and the diagonal elements of $\boldsymbol{\Sigma}$ are real and non-negative. These diagonal elements are the singular values, and the columns of $\mathbf{U}$ and $\mathbf{V}$ are the left singular vectors and right singular vectors respectively.

The SVD provides an intuitive way to understand the action of the matrix $\mathbf{A}$ on any vector $\boldsymbol{q}$. First, $\boldsymbol{q}$ is decomposed in the orthonormal basis provided by the columns of $\mathbf{V}$. Each component is then scaled by the corresponding singular value. The results then give the amplitudes of the response in the basis provided by the columns of $\mathbf{U}$. Hence, the singular values of the transfer and inductance matrix represent the efficiency by which quantities on the control surface (either B_n or κ respectively) propagate to the plasma surface. The left- and right-singular values associated with each singular value represent the modes of B_n or κ which propagate with the given efficiency. The singular vectors of the transfer and inductance matrices give our efficiency-ordered magnetic field distributions.

Let $\boldsymbol{\Phi}^{pl}$ and $\boldsymbol{\Phi}^{fix}$ denote the flux vectors associated with the normal magnetic fields on the plasma surface B_n^{pl} and B_n^{fix} defined in the introduction. If the number of basis functions retained on the plasma and control surfaces is the same, then the singular-valued part of the current potential required to achieve total $B_n \approx 0$ can be computed in principle from $\boldsymbol{I} = \mathbf{M}^{-1}[\boldsymbol{\Phi}^{pl} + \boldsymbol{\Phi}^{fix}]$. We say 'in principle' because $\mathbf{M}$ is very poorly conditioned, for a large current in a sufficiently small loop on the control surface will cause little change to the plasma surface B_n. If the number of basis functions on the surfaces is not the same, one can still apply aforementioned relation if $\mathbf{M}^{-1}$ is interpreted as the Moore-Penrose pseudoinverse, which is equivalent to solving for $\boldsymbol{I}$ in a least-squares sense. This procedure for computing a current potential is essentially the one used in the `NESCOIL` code. Using the SVD of $\mathbf{M} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T$ we can write

$$\boldsymbol{I} = \mathbf{V}\boldsymbol{\Sigma}^{-1}\mathbf{U}^T[\boldsymbol{\Phi}^{pl} + \boldsymbol{\Phi}^{fix}]. \tag{23}$$

To understand why large shaping currents $\boldsymbol{I}$ result from this procedure for some plasma shapes but not for other shapes, we will consider various subsets of the products on the right-hand side of (23). Let $\boldsymbol{S}_e$ denote the rightmost product, $\boldsymbol{S}_e = \mathbf{U}^T[\boldsymbol{\Phi}^{pl} + \boldsymbol{\Phi}^{fix}]$, which is the projection of the externally generated plasma surface B_n onto the left singular vectors (which represent B_n distributions on the plasma surface.) We call the elements of $\boldsymbol{S}_e$ the efficiency sequence. The elements of this sequence are relatively insensitive to the distance between the plasma and control surfaces, since this separation mainly affects the singular values rather than the unitary matrix $\mathbf{U}$. (Recall the eigenvalues of a unitary matrix all have unit norm.) Including the next term in the product (23), we define the feasibility sequence $\boldsymbol{S}_f = \boldsymbol{\Sigma}^{-1}\mathbf{U}^T[\boldsymbol{\Phi}^{pl} + \boldsymbol{\Phi}^{fix}]$, which is the efficiency sequence divided by the singular values. Roughly speaking, the sum of the feasibility sequence determines the magnitude of the required shaping currents $\boldsymbol{I}$, since $\boldsymbol{I} = \mathbf{V}\boldsymbol{S}_f$ and $\mathbf{V}$ is unitary. An increasing feasibility sequence has a diverging sum, so the magnitude of the required shaping currents diverges, hence it is likely infeasible to find suitable coils. Conversely, a rapidly decreasing feasibility sequence means the amplitudes of shaping currents are small, hence it should be feasible to find suitable coils. As the feasibility sequence depends on the singular values, it will be sensitive to the separation between plasma and control surfaces. Efficiency and feasibility sequences can be defined for the transfer matrix just as for the inductance matrix.

IV. NUMERICAL SOLUTION

The inductance matrix can be computed from the definition (20), discretizing the two surface integrals using uniform grids in θ and ζ on each surface. A convenient expression for numerical evaluation of the inductance matrix is

$$\mathbf{M} = (\mathbf{F}^p)^T \mathbf{D} \mathbf{F}^c, \tag{24}$$

where the matrix element $F^p_{i,j}$ is given by basis function j on the plasma surface evaluated at a (θ, ζ) grid point indexed by i, an analogous definition holds for $\mathbf{F}^c$ (representing the basis functions on the control surface),

$$D_{i,j} = \Delta\theta\ \Delta\zeta\ \Delta\theta'\ \Delta\zeta'\ n_p \frac{\mu_0}{4\pi} \sum_{\ell'} \left[\frac{\boldsymbol{N}_i \cdot \boldsymbol{N}'_j}{\left|\boldsymbol{r}_i - \boldsymbol{r}'_j\right|^3} - \frac{3(\boldsymbol{r}_i - \boldsymbol{r}'_j)\cdot \boldsymbol{N}_i (\boldsymbol{r}_i - \boldsymbol{r}'_j)\cdot \boldsymbol{N}'_j}{\left|\boldsymbol{r}_i - \boldsymbol{r}'_j\right|^5} \right], \tag{25}$$

$\Delta\theta$ and $\Delta\zeta$ denote the spacing of grid points, $\boldsymbol{r}_i$ and $\boldsymbol{N}_i$ are evaluated at a (θ, ζ) grid point on the plasma surface indexed by i, and $\boldsymbol{r}'_j$ and $\boldsymbol{N}'_j$ are evaluated at a (θ', ζ') grid point on the control surface indexed by j. Corresponding to (11), the ζ and ζ' sums represent integrals over the interval $[0, 2\pi/n_p]$ (i.e. over one of the identical toroidal periods) and ℓ' indexes the toroidal periods. Due to symmetry the corresponding ℓ sum is reduced to the factor of n_p in (25). Computation times are typically dominated by the assembly of $\mathbf{D}$ and by the matrix multiplications in (24). Calculations can be significantly accelerated with little programming effort by using OpenMP parallelism for the former and a multi-threaded BLAS subroutine for the latter.

The basic resolution parameters of the numerical calculation are the number of grid points N_θ and N_ζ on each surface, determining both dimensions of $\mathbf{D}$ and determining the number of rows of $\mathbf{F}^p$ and $\mathbf{F}^c$, and the maximum poloidal and toroidal mode numbers $m_{\max}$ and $n_{\max}$ for the set of basis functions on each surface, determining the number of columns of $\mathbf{F}^p$ and $\mathbf{F}^c$ and determining both dimensions of $\mathbf{M}$. Typically it is appropriate to choose $N_\theta \gg m_{\max}$ and $N_\zeta \gg n_{\max}$. In principle, different numbers of grid points could be used on the two surfaces, and/or different numbers of basis functions could be used on the two surfaces. However for all results presented here, the resolution parameters are identical on the two surfaces.

Now consider computation of the transfer matrix, which amounts to solving Laplace's equation in the region between the plasma and control surfaces. This problem is a classic case in which boundary integral equation methods excel [11–15], since these methods avoid the need to discretize the volume inside the control surface; only the surfaces need

to be discretized. Within boundary integral equation methods, there are many options. Methods can be divided in two classes, those which use Green's theorem as in [13], and those which use a surface source distribution [11] in which we write a surface integral

$$\phi(\boldsymbol{r}) = \int \mathrm{d}^2 r' \; y(\boldsymbol{r}') G(\boldsymbol{r} - \boldsymbol{r}'). \tag{26}$$

Here, y is some surface distribution which is solved for in terms of the Neumann boundary condition on the control surface, and $G(\boldsymbol{d})$ is a solution of Laplace's equation with a singularity at $\boldsymbol{d} = 0$, such as a monopole or dipole. The sources are outside the control volume (by a finite or infinitesimal distance) and are not physical, introduced to span the space of possible solutions to Laplace's equation inside the control surface, so G can be the monopole function $1/|\boldsymbol{d}|$ even for this magnetic field calculation. For surface source distribution methods, one can choose the surface on which the sources are located (i.e. the integration surface in (26)) to either be infinitesimally outside the control surface, or it can be a distinct third surface which is outside the control surface by a finite distance. In the latter case, the transfer matrix must be independent of the specific choice of outermost surface if the calculation is well resolved numerically, since the transfer matrix is originally defined without reference to this third surface. For good numerical convergence, this outer surface must be far from the control surface compared to the spacing of integration points, but not so far that the coupling between outer and control surfaces is poorly conditioned. The 3-surface method has the disadvantage that selecting a third surface shape which is an appropriate distance from the control surface sometimes requires manual adjustments; however the 2-surface source distribution method and Green's theorem method have the disadvantage that one must carefully integrate over a singularity at the point $\boldsymbol{r}' = \boldsymbol{r}$, as in (26). These singularities may be handled using specialized quadrature rules, or using the remarkable regularization technique of Merkel [13, 14] in which one adds and subtracts a singular function that can be integrated analytically. (Note that the appendix of [13] contains many typographic errors. Readers should refer instead to [14] which contains the correct equations.)

For any of these schemes, the transfer matrix is computed using a similar sequence of steps. First, a dense matrix is assembled which relates either $\phi(\boldsymbol{r})$ (in the Green's theorem method) or $y(\boldsymbol{r})$ (in the surface source method) on the outermost surface to B_n on the control surface. The inverse of this matrix is left-multiplied by another matrix that relates $\phi(\boldsymbol{r})$ or $y(\boldsymbol{r})$ on the outermost surface to B_n on the plasma surface. For example, the method described in [8] for computing the transfer matrix is a 3-surface source distribution method with dipole sources. One computes the outer-to-plasma-surface and outer-to-control-surface inductance matrices $\mathbf{M}^{(p,o)}$ and $\mathbf{M}^{(c,o)}$, and then the transfer matrix follows from $\mathbf{T} = \mathbf{M}^{(p,o)}(\mathbf{M}^{(c,o)})^{-1}$. This method amounts to spanning the set of solutions to Laplace's equation using dipoles on the outer surface, where the dipole amplitudes are chosen to satisfy the Neumann boundary condition on the control surface.

For results shown here, we use two different surface source distribution methods, both the 2-surface scheme with monopole source distribution and Merkel's regularization, and the 3-surface scheme with dipole source distribution, and we have verified the two approaches give identical results as a check of correctness. For the 2-surface method, we use the same uniform (θ, ζ) grid for y and B_n. For every calculation shown in the figures below, computation of the transfer and inductance matrices and their SVDs took in total ≤ 160 s on one node (24 cores) of the NERSC Edison computer, and computations could be accelerated with further code optimization if desired.

V. CIRCULAR AXISYMMETRIC SURFACES

To gain familiarity with the transfer and inductance matrices and their SVDs, let us first consider the case in which both the plasma and control surfaces are axisymmetric, though the magnetic field distributions we consider may be nonaxisymmetric. We begin by analyzing the cylindrical limit, corresponding to large aspect ratio and circular cross section of the surfaces. In this limit, we can compute the singular values and vectors of the transfer and inductance matrices analytically, because the singular vectors correspond to separable solutions of Laplace's equation for the scalar magnetic potential ϕ in cylindrical geometry. This correspondence can be understood as follows. Laplace's equation becomes

$$\frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial \phi}{\partial r}\right) + \frac{1}{r^2}\frac{\partial^2 \phi}{\partial \theta^2} + \frac{1}{R^2}\frac{\partial^2 \phi}{\partial \zeta^2} = 0 \tag{27}$$

where $R \approx$ constant. Some of the separable solutions are

$$\phi = y^s_{m,n} \sin(m\theta - n\zeta) I_m\left(\frac{nr}{R}\right) \quad \text{and} \quad y^c_{m,n} \cos(m\theta - n\zeta) I_m\left(\frac{nr}{R}\right) \tag{28}$$

for integers m and $n \neq 0$, where I_m is the modified Bessel function, and $y^s_{m,n}$ and $y^c_{m,n}$ are arbitrary constants. There are also "axisymmetric" solutions

$$\phi = y^s_{m,0} \sin(m\theta) r^m \text{ and } y^c_{m,0} \cos(m\theta) r^m. \tag{29}$$

Forming $B_n = \partial\phi/\partial r$ and noting the weight w is $\approx$ constant, these fields represent orthogonal flux vectors Φ_j on both the plasma and control surfaces. Thus, a singular value decomposition $\mathbf{U\Sigma V}^T$ for the transfer matrix can be constructed as follows: the columns of $\mathbf{U}$ represent the (normalized) distributions (28)-(29) for various (m,n) evaluated on the plasma surface, the columns of $\mathbf{V}$ represent the (normalized) distributions (28)-(29) for various (m,n) evaluated on the control surface, and the singular values are the ratios of fluxes Φ_j between the plasma and control surfaces. Since the SVD is unique up to the signs of the singular vectors under suitable assumptions (singular values in decreasing order, matrix is square, etc.), then the SVD constructed above must be "the" SVD of the transfer matrix. The singular values $\sigma_{m,n}$ are thus

$$\sigma_{m,n} = \frac{\Phi_P}{\Phi_C} = \frac{a_P}{a_C}\left[\frac{I_{m-1}\left(\frac{na_P}{R}\right) + I_{m+1}\left(\frac{na_P}{R}\right)}{I_{m-1}\left(\frac{na_C}{R}\right) + I_{m+1}\left(\frac{na_C}{R}\right)}\right] \tag{30}$$

for the $n \neq 0$ fields (28), where a_P and a_C are the minor radii of the plasma and control surfaces, and

$$\sigma_{m,0} = (a_P/a_C)^m \tag{31}$$

for the $n = 0$ fields (29). One can see from these formulae that in the limit that the plasma and control surfaces approach each other, $a_P/a_C \to 1$, the singular values become 1.

If n is not as large as the aspect ratio of the surfaces, the Bessel functions in (30) can be expanded for small argument, giving $\sigma_{m,n} \approx (a_P/a_C)^m$ for $m \geq 1$, or $\sigma_{m,n} \approx (a_P/a_C)^2$ for $m = 0$. In the latter case, it is interesting that the singular value becomes independent of the physical wavelength $2\pi R/n$, so the $m = 0$ distributions are actually less efficient (i.e. have singular values that are smaller by a_P/a_C) than the shorter wavelength $m = 1$ distributions. The $m = 0$ distributions are comparable in efficiency to the $m = 2$ distributions.

A similar analysis can be done in this cylindrical limit to compute the singular values of the inductance matrix. In this case, we must also consider the Laplace equation solutions which extend outside the control surface to infinity, replacing $r^m \to r^{-m}$ in (29) and $I_m \to K_m$ (the modified Bessel function of the second kind) in (28). Continuity of B_r at the control surface is imposed, as is a jump condition on the tangential magnetic field associated with the surface current. Supposing the current potential has the form

$$\kappa(\theta,\zeta) = \kappa^s_{m,n} \sin(m\theta - n\zeta) + \kappa^c_{m,n} \cos(m\theta - n\zeta) \tag{32}$$

for some $\kappa^s_{m,n}, \kappa^c_{m,n}$, we thereby obtain the radial magnetic field inside the control surface to be

$$\begin{aligned} B_r = &-\frac{\mu_0 n}{2R}\left[\frac{I_{m-1}\left(\frac{na_C}{R}\right) + I_{m+1}\left(\frac{na_C}{R}\right)}{K_{m-1}\left(\frac{na_C}{R}\right) + K_{m+1}\left(\frac{na_C}{R}\right)} K_m\left(\frac{na_C}{R}\right) + I_m\left(\frac{na_C}{R}\right)\right]^{-1} \left[I_{m-1}\left(\frac{nr}{R}\right) + I_{m+1}\left(\frac{nr}{R}\right)\right] \\ &\times \left[\kappa^s_{m,n} \sin(m\theta - n\zeta) + \kappa^c_{m,n} \cos(m\theta - n\zeta)\right] \end{aligned} \tag{33}$$

when $n \neq 0$, and

$$B_r = -\frac{\mu_0 m}{2} \frac{r^{m-1}}{a_C^m} \left[\kappa^s_{m,0} \sin(m\theta) + \kappa^c_{m,0} \cos(m\theta)\right] \tag{34}$$

when $n = 0$. Computing the magnetic flux Φ_P on the plasma surface and the current I_C on the control surface, the singular values $\sigma_{m,n}$ then follow from their ratio:

$$\begin{aligned} \sigma_{m,n} &= \left|\frac{\Phi_P}{I_C}\right| \\ &= 2\pi^2 n \mu_0 a_P \left[\frac{I_{m-1}\left(\frac{na_C}{R}\right) + I_{m+1}\left(\frac{na_C}{R}\right)}{K_{m-1}\left(\frac{na_C}{R}\right) + K_{m+1}\left(\frac{na_C}{R}\right)} K_m\left(\frac{na_C}{R}\right) + I_m\left(\frac{na_C}{R}\right)\right]^{-1} \left[I_{m-1}\left(\frac{na_P}{R}\right) + I_{m+1}\left(\frac{na_P}{R}\right)\right] \end{aligned} \tag{35}$$

when $n \neq 0$, and

$$\sigma_{m,0} = 2\pi^2 m \mu_0 R \left(a_P/a_C\right)^m \tag{36}$$

for $n = 0$. These formulae, unlike those for the transfer matrix, remain dependent on m, n, and the aspect ratio even when the plasma and control surfaces coincide, $a_P/a_C \to 1$.

Some insight into (35) can be gained by expanding the Bessel functions for small argument, which is valid when n is below the aspect ratio. In this case, when m is nonzero, we obtain again (36), indicating the singular values are nearly independent of n. However, when $m = 0$, the expansion of (35) yields instead

$$\sigma_{0,n} \approx 2\pi^2 n^2 \mu_0 a_P^2 / R. \tag{37}$$

This expression indicates the singular values actually grow like n^2, so the dominant singular vectors correspond to short wavelengths rather than long wavelengths. This behavior is one instance of a phenomenon mentioned in the introduction: the inductance matrix sometimes treats long-wavelength field distributions as inefficient, because the field is driven by the *gradient* of the current potential, and this gradient is small for long wavelengths. At sufficiently large n however the approximation leading to (37) breaks down, and the singular values begin to decrease with n.

For the geometries in figure 1, the singular values of both transfer and inductance matrices computed using these approximate analytic formulae are shown in figure 2. The singular values of the inductance matrix have been normalized to the characteristic value $2\pi^2\mu_0 R$ so the largest singular values are $O(1)$. Overall, the distributions of singular values for the transfer and inductance matrices are quite similar, aside from two factors. First, the $m = 0$ behavior is different, as already discussed, with the singular values of the transfer matrix decreasing monotonically with n whereas those of the inductance matrix peak at intermediate n. Second, the singular values of the transfer matrix generally decrease faster with m and n compared to the singular values of the inductance matrix, and so a different color scale is used for the two matrices in the figure. The physical reason for this different rate of decrease is the same effect described in the previous paragraph.

One practical consequence of the analysis in this section is that we can analytically determine the toroidal mode numbers which can be controlled with comparable efficiency to given poloidal mode numbers. In other words, given some poloidal mode number m_0, what is the toroidal mode number n_0 such that the $(m = m_0, n = 0)$ and $(m = 0, n = n_0)$ modes have comparable efficiency, as determined by the singular values? To proceed, we compute $\sigma_{m=0,n=n_0}$ from (30), expanding the Bessel functions for large argument, $I_m(x) \approx e^x/\sqrt{2\pi x}$. Equating the result to $\sigma_{m=m_0,n=0}$ from (31), we find

$$n_0 = \left(m_0 - \frac{1}{2}\right) \frac{R}{a_P} \left(\frac{a_C}{a_P} - 1\right)^{-1} \ln\left(\frac{a_C}{a_P}\right). \tag{38}$$

In figure 2 we show ellipses in (m, n)-space, where the extent in n is determined by (38) for several choices of m_0. It can be seen that (38) indeed gives quite an accurate estimate for the range of m and n which have comparable efficiency. While (38) was derived based on the transfer matrix singular values, it does not appear possible to form such an explicit expression for the inductance matrix singular values due to their more complicated dependence on m and n. However the transfer and inductance matrix contours in figure 2 are quite similar.

Next, we compare these analytic results to direct numerical calculations which fully account for toroidal effects. First considering only axisymmetric basis functions for simplicity, results are shown in figure 3 for three values of the minor radius of the plasma surface, taking the aspect ratio of the control surface to be $R/a_C = 10$. The agreement with (31) and (36) is extremely close. The B_n distributions given by the singular vectors are found to correspond nearly exactly to Fourier modes of increasing m, as expected, coming in pairs with $\sin(m\theta)$ or $\cos(m\theta)$ symmetry. The factor of m in (36) which is not present in (31) causes the singular values of the inductance matrix to decrease less rapidly than those of the transfer matrix, particularly at low n.

Next, we consider the set of basis functions given by various n for $m = 0$, with results shown in figure 4. Excellent agreement is seen with the analytical results (30) and (35). As expected, the B_n distributions given by the singular vectors are found to correspond nearly exactly to Fourier modes of increasing n for the transfer matrix, and to Fourier modes of *decreasing* n for the inductance matrix. For these figures, numerical parameters used were $N_\theta = 64$ and $N_\zeta = 256$.

Figure 5 mirrors figure 2, but showing data from the fully toroidal numerical calculation. In this calculation each singular vector contains a mixture of m and n values, so the figure shows the singular value associated with the left singular vector that has maximal amplitude of each given (m, n). Figures 2 and 5 are not exactly identical, due to toroidal effects which are not accounted for in the analytic expressions plotted in the former figure. However, the agreement is generally quite good, demonstrating that the analytic expressions (30), (31), (35), and (36) are remarkably accurate even at realistic values of aspect ratio. The predicted non-monotonic n dependence of the $m = 0$ singular values of the inductance matrix can be clearly seen.

Equations (30)-(31) and (35)-(36) may actually be useful in applications that do not directly involve the transfer or inductance matrices, since they provide an efficiency ordering to pairs of poloidal and toroidal mode numbers (m, n).

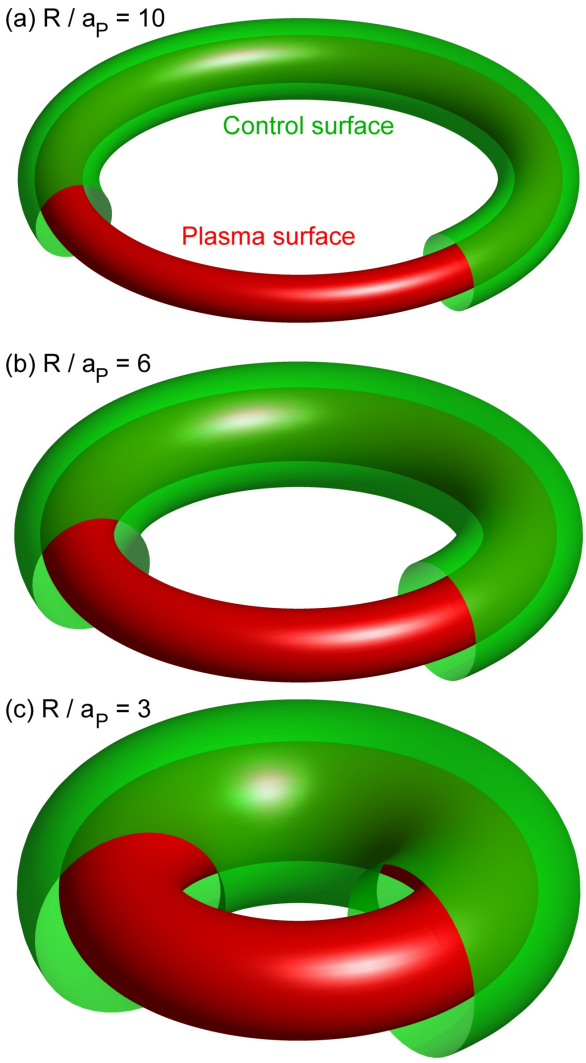


FIG. 1. (Color online) Geometries used for figures 2 and 5. The ratio a_C/a_P (minor radius of the control surface divided by that of the plasma surface) is 1.7 in all three cases.

For example, considering the Dommaschk potentials [16] for vacuum magnetic fields, the potentials for various (m, n) numbers can be sorted by $\sigma_{m,n}$ to order them by efficiency.

Finally, figure 6 illustrates the behavior of the transfer matrix SVD at finite aspect ratio. Figure 6.a shows the geometry, with the plasma aspect ratio R/a_P taken to be 3, and a_C/a_P chosen to be 1.7, representing a reasonable distance from the plasma at which control coils might be placed. Considering both sine and cosine symmetries, figure 6.b shows there are approximately 500 distributions of B_n with singular values above 0.01. For comparison, the results of (30)-(31) are also plotted, although the large-aspect-ratio approximation used to derive these analytic results is not very well satisfied. Despite this questionable approximation, the analytic method does a good job of predicting the distribution of the 'true' (numerical) singular values. Numerical parameters used were $N_\theta = 64$, $N_\zeta = 256$, $m_{\max} = 24$, and $n_{\max} = 60$. We verified that none of the points in figure 6 moved visibly when any of the resolution parameters was doubled.

Red crosses in figure 6.b show the numerically computed singular values if only the axisymmetric distributions are retained. These singular values correspond to translation, elongation, triangularity, squareness, etc., and they come in pairs (barely visible in the plot) corresponding to sine and cosine symmetry. It can be clearly seen that to control B_n to any given tolerance, there are many more nonaxisymmetric degrees of freedom than there are axisymmetric degrees of freedom.

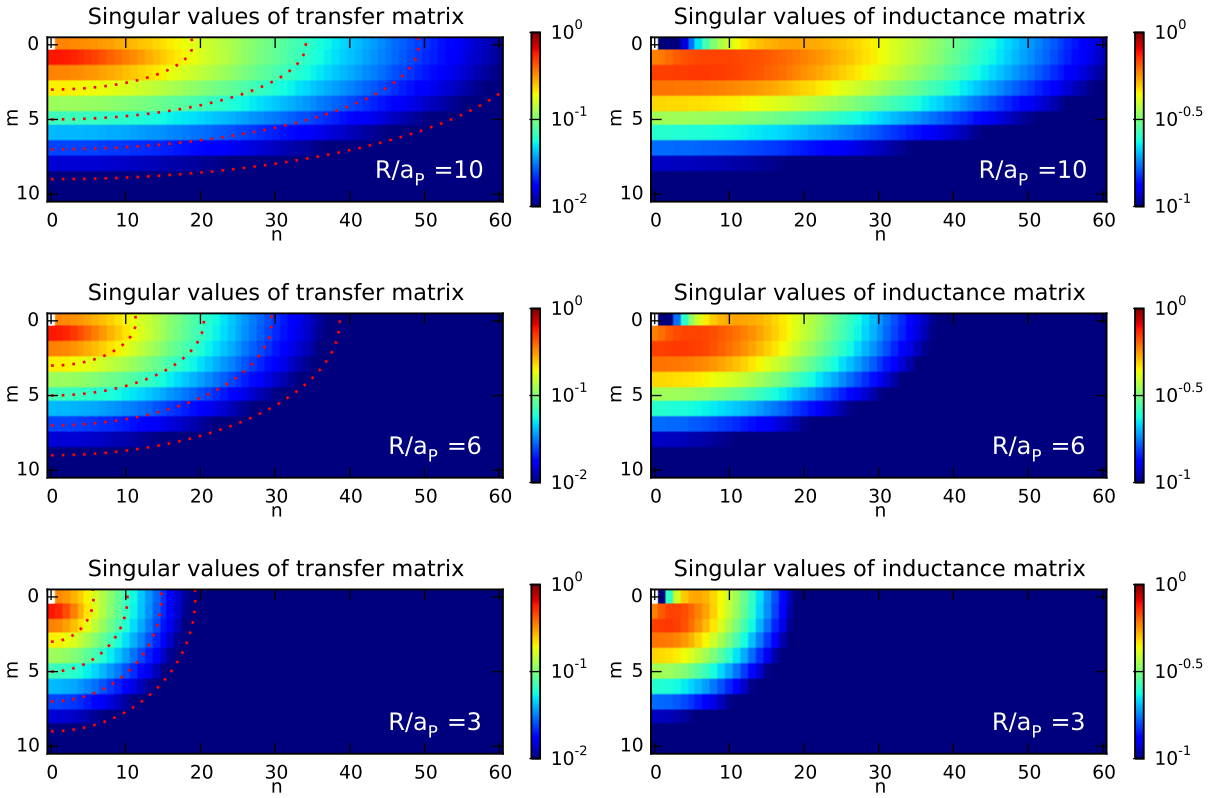


FIG. 2. (Color online) Singular values of the transfer and inductance matrices, computed using the approximate analytic expressions (30), (31), (35), and (36), valid for large aspect ratio and circular cross section. Results are shown for three values of aspect ratio, using $a_C/a_P = 1.7$. The inductance matrix singular values have been normalized by $2\pi^2\mu_0 R$. Note the color scale is different for the transfer and inductance matrices. Red dashed curves are ellipses with major (minor) axes n_0 (m_0) related by (38).

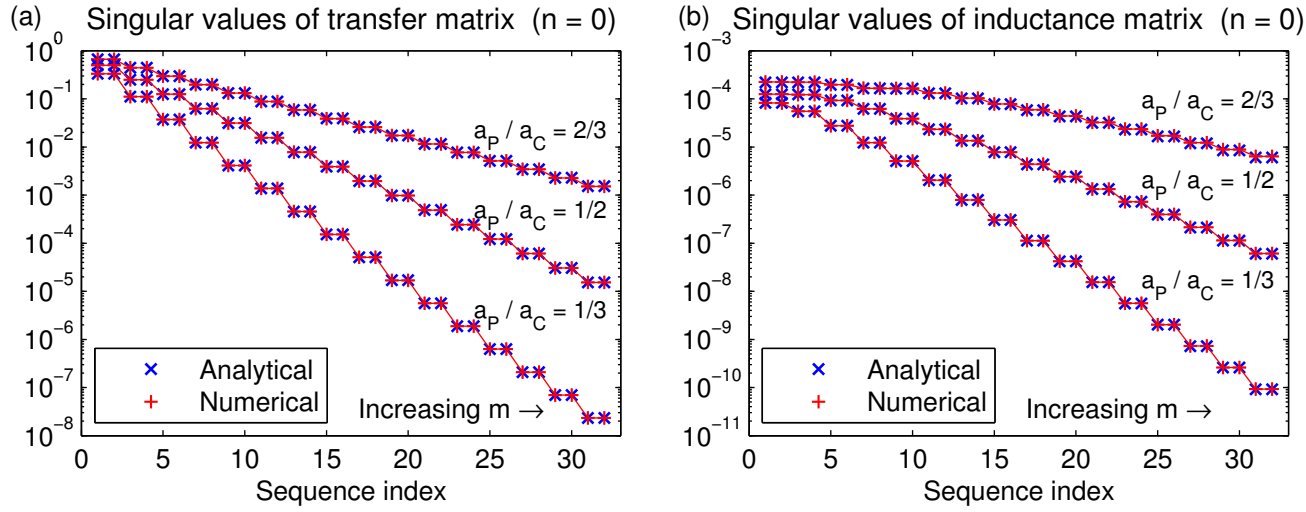


FIG. 3. (Color online) Singular values of the (a) transfer matrix and (b) inductance matrix for surfaces of circular cross-section at high aspect ratio, considering only axisymmetric basis functions. Fully toroidal numerical calculations agree extremely well with the predictions (31) and (36) from the separable cylindrical solutions to Laplace's equation. The singular values appear in pairs because both sine and cosine phases are permitted.

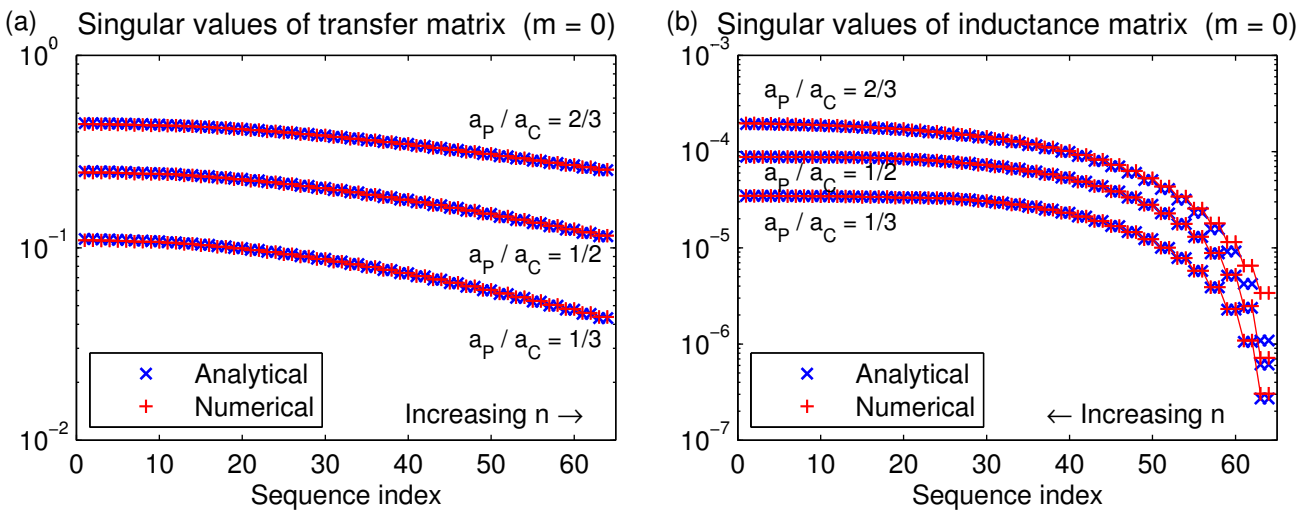


FIG. 4. (Color online) Singular values of the (a) transfer matrix and (b) inductance matrix for surfaces of circular cross-section at high aspect ratio, considering only poloidally constant basis functions. Fully toroidal numerical calculations agree extremely well with the predictions (30) and (35) from the separable cylindrical solutions to Laplace's equation.

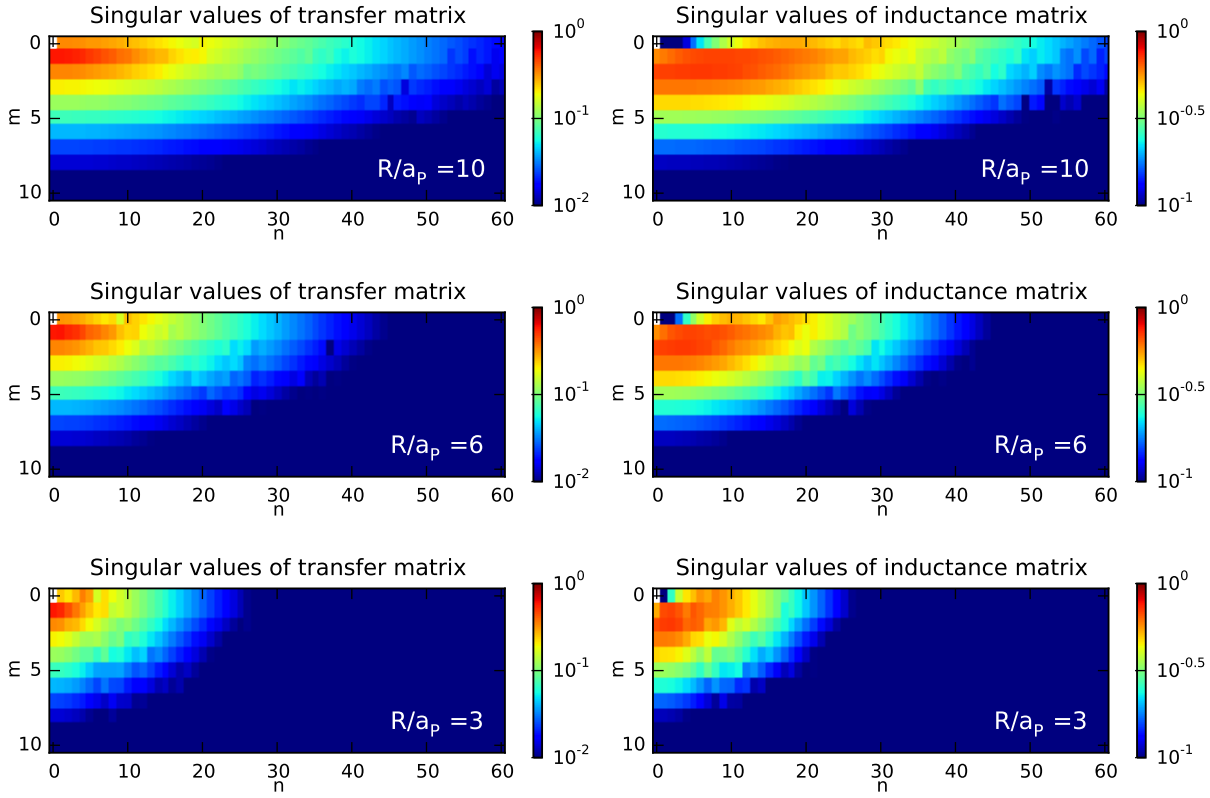


FIG. 5. (Color online) Singular values of the transfer and inductance matrices, computed using direct numerical simulation and fully accounting for toroidal effects, sorted by dominant poloidal and toroidal mode numbers (m, n). Results are shown for three values of aspect ratio, using $a_C/a_P = 1.7$. The inductance matrix singular values have been normalized by $2\pi^2\mu_0 R$. Note the color scale is different for the transfer and inductance matrices. The similarity to figure 2 demonstrates the good accuracy of the analytic formulae even at finite aspect ratio.

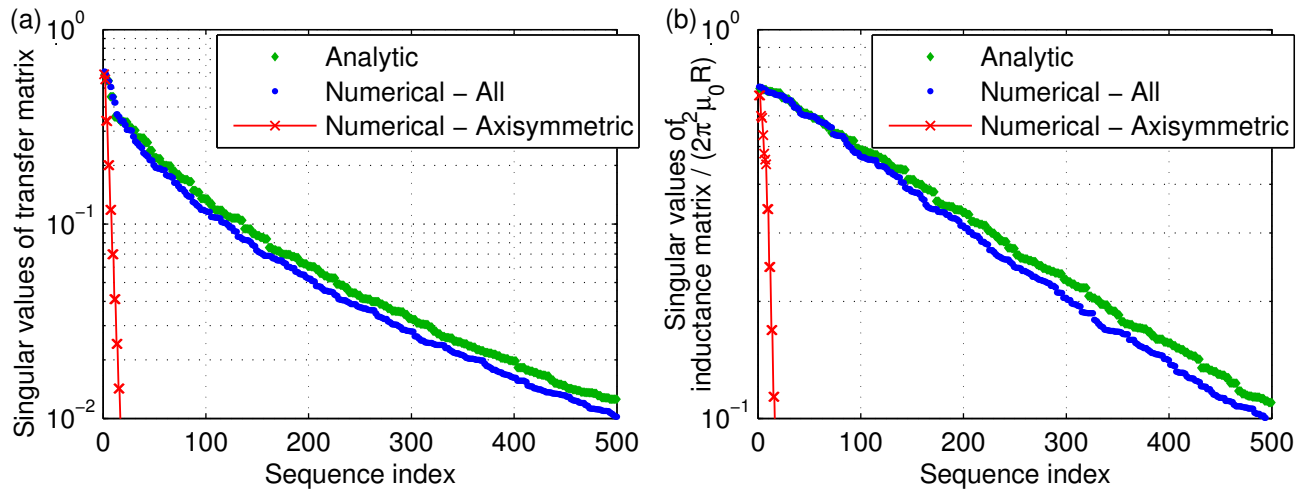


FIG. 6. (Color online) Singular values for the geometry $R/a_P = 3$ and $a_C/a_P = 1.7$ (figure 1.c). Values computed by the code are shown as blue circles. The subset of these singular values associated with axisymmetric singular vectors are shown as red crosses. The cylindrical analytic results (30)-(31) and (35)-(36) are shown as green diamonds for comparison.

VI. SHAPED AXISYMMETRIC SURFACES

Now that the singular vectors and values of the inductance and transfer matrices are known for the case of circular surfaces, we proceed to consider plasma and control surfaces which are axisymmetric but non-circular. At this point we will be able to demonstrate one of our main results: both the inductance and transfer matrix methods can correctly identify the X-point divertor shape as being efficient, despite its sharp corner. We will also give examples of plasma shapes that have low curvature and low spectral width, yet are inefficient to produce.

The four shapes we will consider are shown in figure 7. The 'divertor' shape is taken from a numerical Grad-Shafranov solution for the ITER 15 MA baseline scenario. Since our numerical methods presume differentiable surface shapes, the surface at normalized poloidal flux = 0.995 rather than 1 is chosen so the corner at the X-point is very slightly rounded off. The other three shapes are chosen to have small curvature and small spectral width, in the sense that all can be represented by Fourier series with 4 or fewer terms:

$$R(\theta) = \sum_{m=0}^{3} R_m \cos(m\theta), \quad Z(\theta) = \sum_{m=0}^{3} Z_m \sin(m\theta). \tag{39}$$

(For contrast, we retain 100 terms in the Fourier series for the divertor shape.) The 'ellipse' shape is chosen to have the same width, height, and median major radius as the divertor shape. The 'peanut' and 'H' shapes are chosen just large enough to encircle the divertor shape, so any difficulty generating these two shapes cannot be due to larger distance from a common control surface or coils. For the ellipse, peanut, and H shapes, we have generated MHD equilibrium solutions just as for the divertor shape, this time using the VMEC code [10] run in fixed-boundary mode. To emphasize that the divertor shape has a significantly sharper corner than the other shapes, figure 8 shows the curvature (i.e. inverse radius of curvature) of the poloidal cross-section

$$\kappa_{RZ} = (R'Z'' - Z'R'')/(R'^2 + Z'^2)^{3/2}, \tag{40}$$

where primes indicate $d/d\theta$, and κ_{RZ} is independent of the poloidal coordinate θ used.

Intuitively, we expect the ellipse and divertor shapes to be efficient and the peanut and H shapes to be inefficient in some sense, for plasmas of roughly the ellipse and divertor shapes are common in experiments whereas the peanut and H shapes are uncommon. One calculation supporting this notion of efficiency is presented in figure 9. Here, we first compute B_n^{pl}, the magnetic field normal to the desired plasma surface driven by plasma current inside it, using the virtual casing principle [17, 18]. Next, we solve the linear least-squares problem to determine the currents in the actual ITER shaping coils (6 poloidal field coils and 6 central solenoid coils) which minimize $\chi^2 = \int \mathrm{d}^2a\ B_n^2$, where $B_n = B_n^{pl} + B_n^x$ (in this case $B_n^{fix} = 0$) and B_n^x is the contribution from the shaping coils. Contours are then plotted of the total poloidal flux $\Psi(R,Z) = \int_0^R 2\pi R' B_Z(R',Z)\, dR'$, summing the plasma and external contributions. It is clear that the ellipse and divertor shapes can be produced to high accuracy by the ITER coilset, whereas the

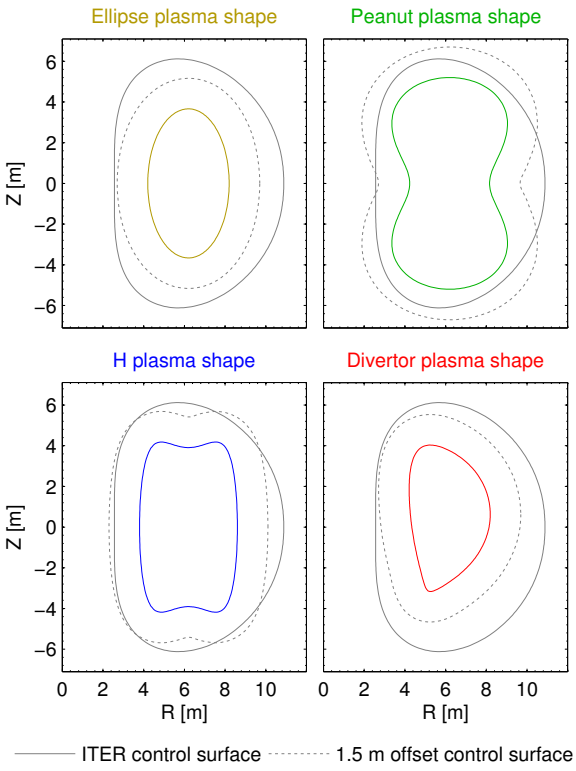


FIG. 7. (Color online) The four noncircular axisymmetric plasma shapes considered in section VI. For each case, the two control surfaces considered are shown in gray.

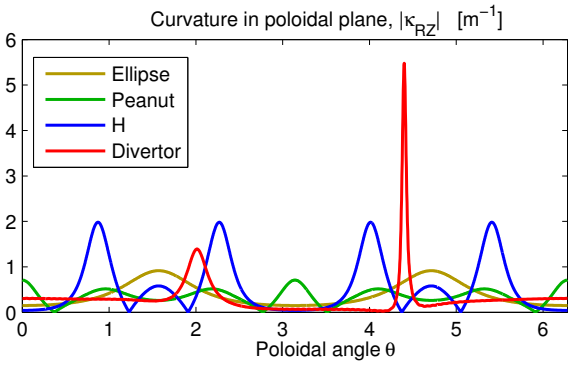


FIG. 8. (Color online) Curvature (= inverse radius of curvature) in the RZ plane for the four plasma surface shapes considered in section VI. For the divertor shape, the spike in curvature corresponds to the (near) X point.

least-squares approximations to the peanut and H shapes lack the desired concave regions. This failure comes despite the fact, mentioned above, that the peanut and H target shapes are everywhere closer to the coils compared to the divertor shape. (The maximum coil currents for these least-squares fits are 29 MA, 13 MA, 45 MA, and 19 MA for the ellipse, peanut, H, and divertor shapes respectively.) If additional coils are included between the actual ITER coils at similar distances from the plasma, the concave plasma shapes can be better approximated, but only with an unrealistic number of coils and large coil currents.

For calculations using the inductance and transfer matrices, we will explore the effect of the choice of control surface by considering two options. One option, the 'ITER control surface', is obtained by fitting the locations of the ITER poloidal field and central solenoid coils with a 4-term Fourier series as in (39). This control surface thus represents a realistic surface on which shaping coils may be placed in an experiment. For contrast, the other option is an 'offset' control surface obtained by expanding the plasma surface by a uniform 1.5 m distance. This choice will enable us to demonstrate that tailoring the control surface to the plasma shape has little effect on our results. The control surfaces

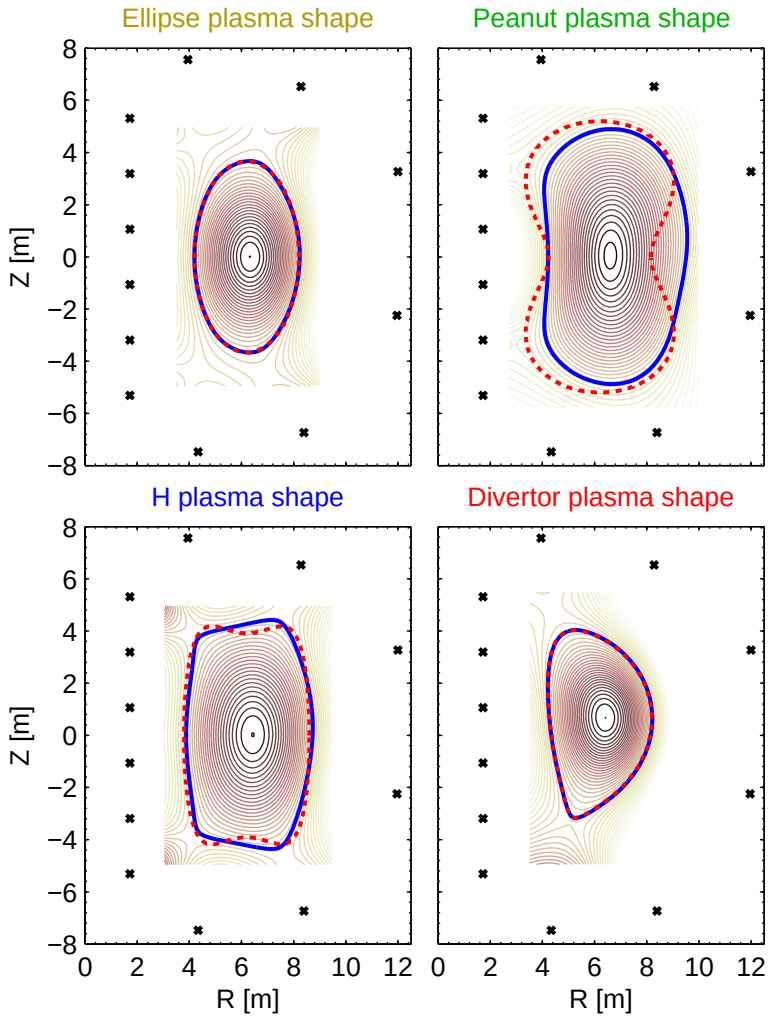


FIG. 9. (Color online) One demonstration that the ellipse and divertor shapes are more efficient to produce, in some sense, than the peanut and H shapes. Red dashed curves show the target plasma shape. Thin solid curves are contours of poloidal flux Ψ when currents in the ITER coilset ($\times$) are chosen so to minimize $\int \mathrm{d}^2 a\, B_n^2$ on the target shape. The achieved Ψ contour closest to the target shape is highlighted with a thick blue curve. While the ellipse and divertor shapes are produced to high accuracy, the desired concave regions in the peanut and H shapes are not produced.

are shown in figure 7. For these calculations we consider only axisymmetric basis functions and singular vectors, and we allow up-down asymmetry.

Our main result is figure 10, which shows the efficiency sequences for the four shapes, i.e. the projection of $B_n^{pl}+B_n^{fix}$ onto the left singular vectors. The toroidal field can be ignored as it gives no B_n, and we can thus take $B_n^{fix}=0$. The four equilibria have the same total plasma current (15 MA) and hence comparable poloidal field, so it is meaningful to compare the absolute magnitudes in figure 10 in addition to the rates of decrease. The figure shows that to cancel B_n^{pl} to a relative accuracy of $\leq 10^{-3}$, only the first 20-25 efficiency-ordered $\boldsymbol{B}$ distributions are required for the ellipse and divertor shapes. The divertor shape is not significantly harder to produce than the ellipse in this regard. In contrast, the 'long tails' for the peanut and H shapes in figure 10 indicate that many more $\boldsymbol{B}$ distributions are required to create these shapes. For the peanut shape, even with the first 80 efficiency-ordered $\boldsymbol{B}$ distributions, it is not possible to cancel B_n^{pl} to a relative accuracy of $< 10^{-2}$. The trends are nearly identical for the transfer and inductance matrices.

Comparing the top row of subplots in figure 10 to the bottom row, it can be seen that changing the control surface may change the precise values of the efficiency sequence, but the overall rate of decrease is not significantly changed. We have also repeated the analysis with other control surfaces, including surfaces with elliptical or circular cross-section, and the results are always quite similar to figure 10. This robustness is expected, for as pointed out previously, the efficiency sequences are independent of the singular values, and hence largely independent of the distance between plasma and control surfaces. Thus, the rates of decrease of the data in figure 10 are a good measure

of the intrinsic efficiency of producing a plasma shape.

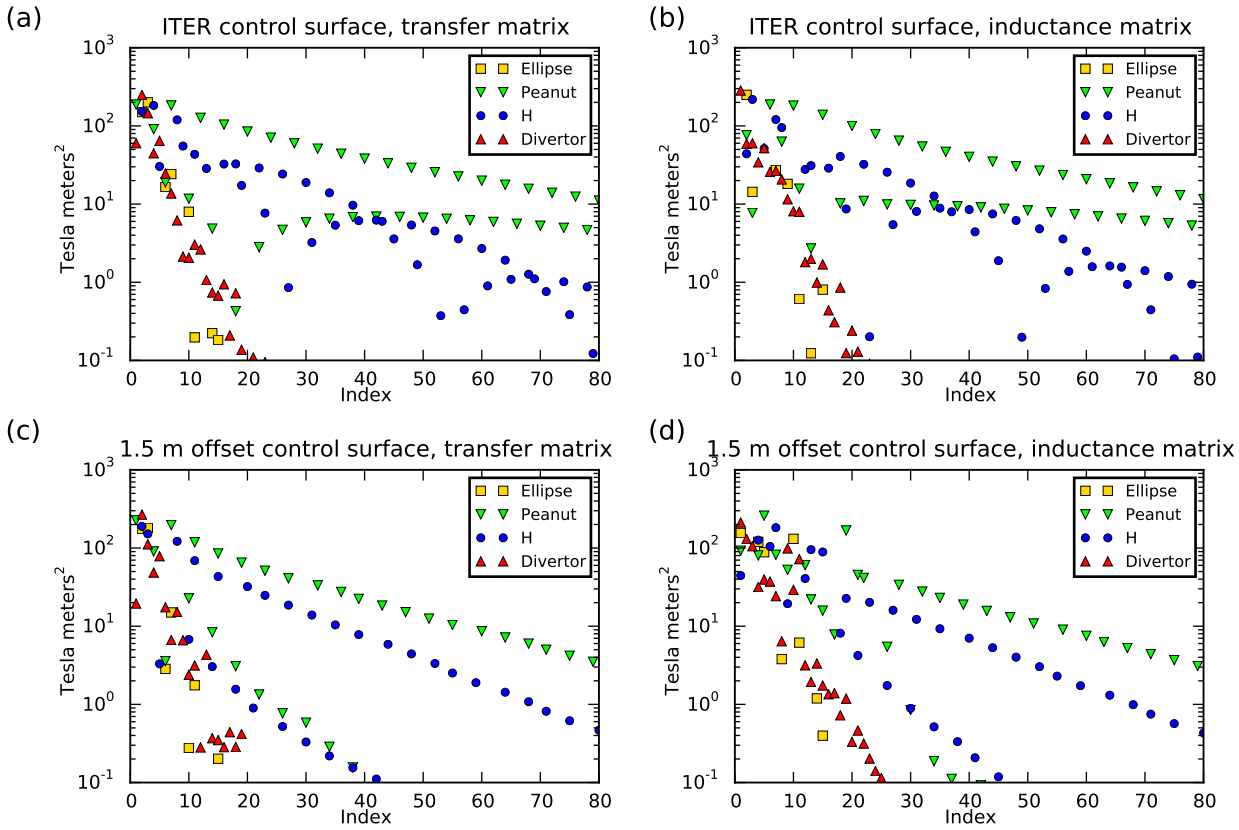


FIG. 10. (Color online) Efficiency sequences, i.e. the projection of B_n^{pl} (magnetic field driven by plasma current, component normal to the target plasma shape) onto the left singular vectors of the inductance and transfer matrices. Comparing the 4 plasma shapes in figure 7, the divertor and ellipse shapes are efficient to produce in the sense that only a small number of singular vectors are required to represent B_n^{pl} to any desired accuracy. Many more singular vectors are required to represent B_n^{pl} for the difficult peanut and H shapes. For the ellipse, peanut, and H shapes, approximately every other sequence element is exactly zero due to up-down symmetry, and the remaining elements alternate in magnitude associated with in-out symmetry vs. asymmetry.

To better understand these results, figure 11 displays additional details of the calculation for the divertor plasma shape with ITER control surface. Figure 11.a shows the field lines of the part of the magnetic field driven by the plasma current. It can be seen that this magnetic field points into the plasma surface on one side of the X point and out of the plasma surface on the other side of the X point. Hence, B_n^{pl} has a sharp transition from one sign to the other in this region (figure 11.b). We use high numerical resolution ($N_\theta = 512, N_\zeta = 96, m_{\max} = 128$) so this transition is well resolved numerically. Figure 11.c then displays the first few right singular vectors of the transfer matrix, corresponding to the magnetic field component normal to the control surface for the most efficient magnetic field distributions. These functions resemble $\cos(\theta)$, $\sin(\theta)$, $\cos(2\theta)$, $\sin(2\theta)$, etc., corresponding to vacuum fields with slow spatial variation, as expected from section V. The first right singular vectors of the inductance matrix (11.d) have slightly faster spatial variation, consistent with the trend we observed in section V. The components of these $\boldsymbol{B}$ distributions normal to the plasma surface are plotted in figures 11.e-f, corresponding to the left singular vectors. These singular vectors have fine structure near $\theta = 4.4$ corresponding to the X point, for exactly the same reason B_n^{pl} does: a $\boldsymbol{B}(\boldsymbol{r})$ with slow variation is projected normal to a surface with a sharp corner. Hence, the basis of left singular vectors is able to resolve fine structure near the X point, precisely where $B_n^{pl}(\theta)$ has fine structure.

Next, in figure 12 we examine the singular values for the four shapes and two control surfaces. The overall pattern of singular values is extremely similar to the case of circular surfaces, figure 3. The singular values for circular surfaces behaved like $(a_P/a_C)^m$ (in pairs with sin or cos symmetry), and we can interpret the singular values for shaped surfaces as following the same pattern for an effective ratio of surface radii a_P/a_C. Particularly for the ellipse and divertor shapes, for which the 1.5 m offset control surface is everywhere closer to the plasma than the ITER control surface, the rate of decrease in the singular values in figure 12 is significantly faster for the more distant surface, as expected. For the ITER control surface, the singular values for the peanut and H shapes are larger than for the

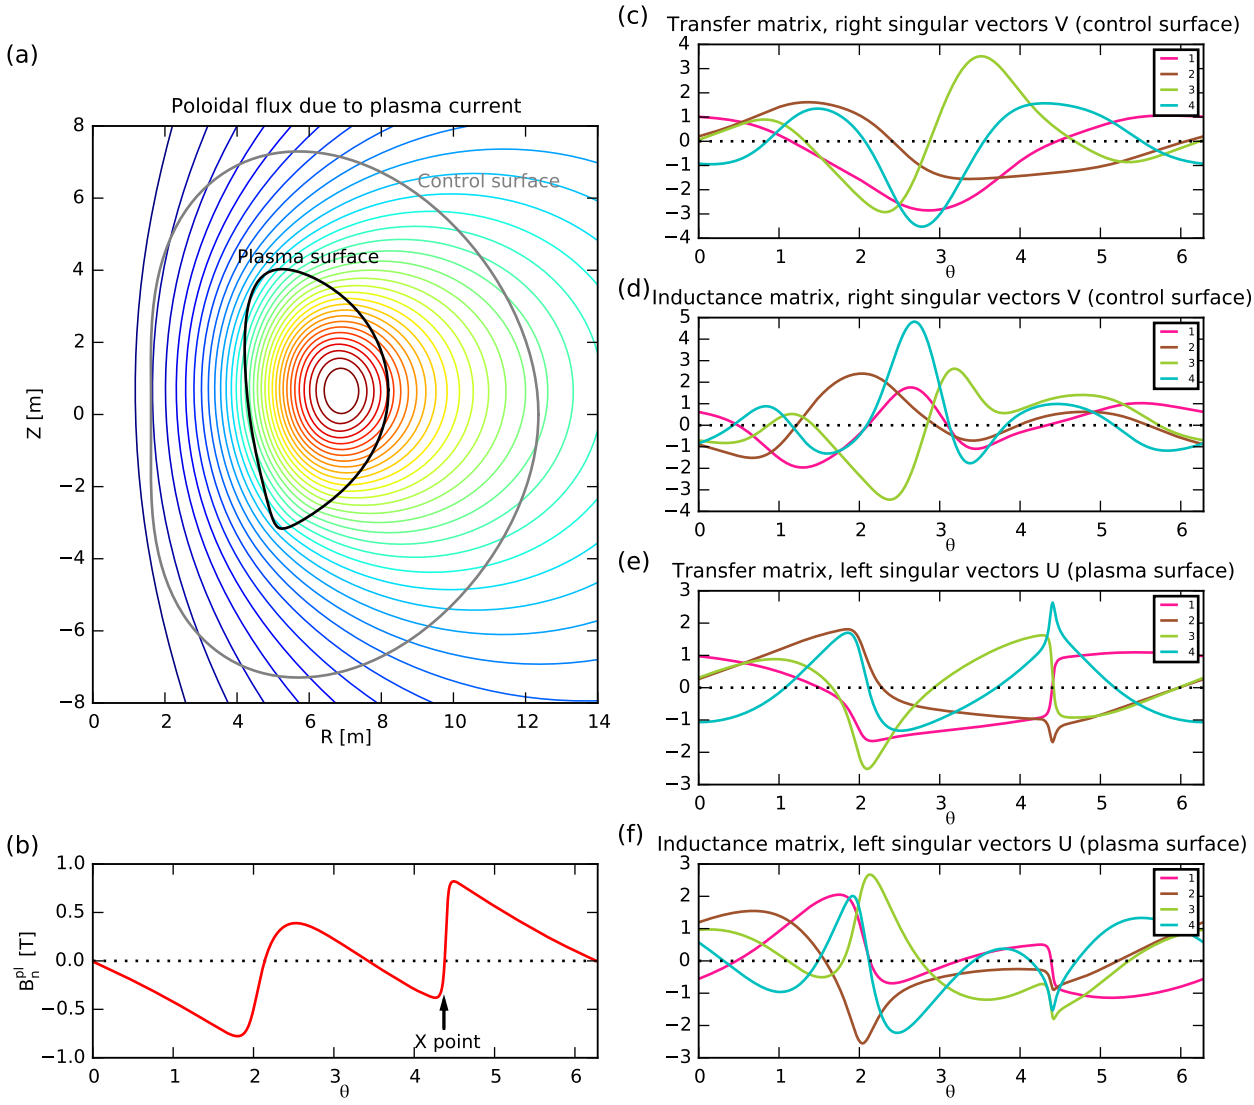


FIG. 11. (Color online) (a) Black and gray curves show the divertor plasma surface and ITER control surface, and colored curves show contours of the poloidal flux Ψ driven by plasma current. These colored contours are equivalently lines of the magnetic field driven by the plasma current. (b) Component of this field normal to the plasma surface. (c) The first 4 right singular vectors of the transfer matrix and (d) inductance matrix are slowly varying, corresponding to vacuum $\boldsymbol{B}$ distributions with slow spatial variation. The left singular vectors correspond to the projection of these distributions normal to the plasma surface, hence they have fine structure near the X point.

ellipse and divertor shapes, which merely reflects the fact that the ellipse and divertor shapes are smaller. For a given effective a_P/a_C, the singular values are not very sensitive to the details of the surface shapes. Hence the singular values themselves are not an ideal target for optimization. The projections shown in figure 10 are more sensitive to the plasma shape, making them a better target for optimization.

Finally, we divide the efficiency sequences (figure 10) by the singular values (figure 12) to get the feasibility sequences in figure 13. Depending on whether the efficiency sequences or singular values decrease more rapidly, the feasibility sequences may be (exponentially) increasing or decreasing. Indeed, the sequence for the ellipse shape is decreasing in the figure, while the sequences for the peanut and H shapes are increasing. The sequences for the divertor shape have a significantly lower slope than the sequences for the peanut and H shapes, reflecting the fact that the divertor shape is easier to generate despite its larger curvature. (We only show the first 24 sequence elements for the divertor shape, since after this point the projection becomes polluted by discretization error from the numerical ITER Grad-Shafranov solution we have available.)

We have repeated the analysis of figures 10 and 13 taking B_n^{pl} to be generated by an infinitesimally thin ring of

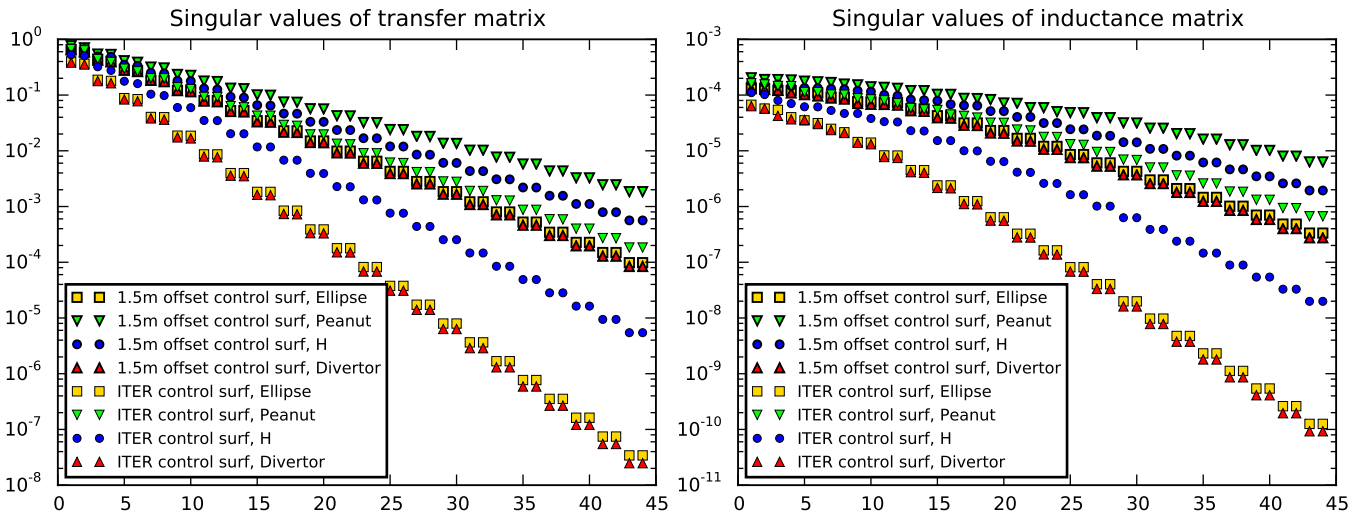


FIG. 12. (Color online) Singular values of the (a) transfer and (b) inductance matrices for the 4 shapes and 2 choices for control surface in figure 7 (considering only axisymmetric distributions.) Despite the strong surface shaping, the results bear close resemblance to the results for circular cross-section surfaces in figure 3.

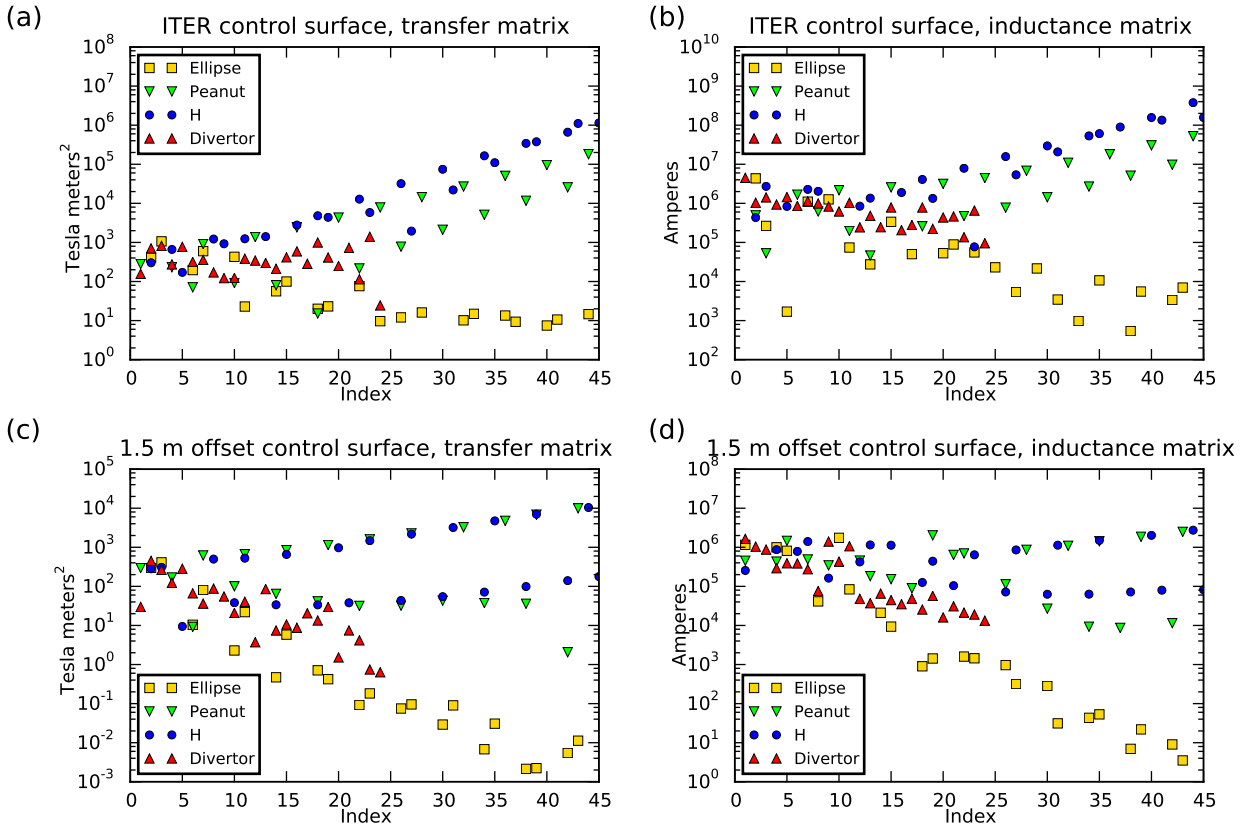


FIG. 13. (Color online) Feasibility sequences, i.e. the efficiency sequences (figure 10) divided by the singular values (figure 12).

current at the magnetic axis, and find only minor changes to the results. The details of the current distribution inside the surface are unlikely to have a strong effect on B_n^{pl} generally since the plasma current is usually concentrated near the magnetic axis, relatively far from the boundary where B_n^{pl} is evaluated.

VII. NONAXISYMMETRIC SURFACES

Now that some understanding of the inductance and transfer matrix SVD properties has been developed in the previous sections, let us present one example from the important but more complicated case of nonaxisymmetric surfaces. Some features from the previous axisymmetric cases will persist. Application of the inductance and transfer matrix SVD for nonaxisymmetric shape optimization will be considered in future work.

Figures 14-18 present the SVD of the transfer and inductance matrices for a W7-X case. The plasma surface is taken from a fixed-boundary `VMEC` equilibrium ($\beta = 4.5\%$) from prior to the selection of the W7-X coil shapes, so there is no ripple due to the discrete coils of the actual experiment. We consider two control surfaces. The 'real' control surface is chosen to be a surface on which the W7-X coils lie. The distance between the plasma and this control surfaces is not exactly uniform. (Not surprisingly, based on the difficulty of generating concave shapes as discussed in the previous section, the designers of W7-X chose to make the coil surface closer to the plasma in the concave regions.) We also consider an 'expanded' control surface, obtained by adding $(0.5\text{m})\cos(\theta)$ to R and $(0.5\text{m})\sin(\theta)$ to Z of the real control surface. For computing the inductance and transfer matrices, in light of the 5-fold toroidal periodicity of W7-X, we enforce 5-fold toroidal periodicity in the basis functions. We also force B_n to have stellarator symmetry by including only stellarator-symmetric basis functions in the calculation.

Figure 14.c shows the singular values of the transfer and inductance matrices. Just as in the circular axisymmetric case of section V, the singular values decrease roughly exponentially, and the rate of decrease is a bit faster for the transfer matrix. The rate of decrease is also faster for the expanded control surface since it is farther from the plasma.

Figures 15-16 show the left and right singular vectors of the transfer matrix for the real control surface, transformed to (θ, ζ)-space. Comparing these two figures, the patterns on the plasma and control surfaces are similar. As one would expect, there is a general trend for the spatial frequencies to become higher as one proceeds through the sequence of singular vectors.

Figures 17-18 show the left and right singular vectors of the inductance matrix for the real control surface, again transformed to (θ, ζ)-space. The spatial frequencies of these singular vectors are noticeably higher than those of the transfer matrix singular vectors, particularly on the control surface (figure 18). This phenomenon is yet another instance of the inductance matrix favoring shorter wavelengths, due to the gradient in the relation between current potential and magnetic field.

Finally, in figure 19 we examine the difficulty of achieving a total $B_n = 0$ on the W7-X plasma surface, showing the efficiency sequences given by projecting B_n^{fix} onto the left singular vectors of the transfer and inductance matrices. Since the surface shapes are nonaxisymmetric, the net poloidal current on the control surface (which we take to flow at constant ζ, as in `NESCOIL`) produces a nonzero B_n^{fix}. (The trends in figure 19 are unchanged if B_n^{fix} is considered to arise instead from an axisymmetric toroidal magnetic field $\propto 1/R$.) The plasma current in W7-X is small, and in this example B_n^{pl} is negligible compared to B_n^{fix}. Figures 19.a-b illustrate that the general rates of decrease in the efficiency sequences are the same for the two control surfaces, as expected. To achieve a total $B_n = 0$ on the plasma surface by adding non-planarity to modular coils on the control surface, the required magnitudes of the B_n or current distributions on the control surface are obtained by dividing the figure 19.a-b sequences by the singular values, yielding the feasibility sequences. Results are shown in figures 19.c-d. For the real control surface, the rate of decrease in figure 19.a-b is just barely faster than the rate of decrease in the singular values, so the resulting sequences in figures 19.c-d show a trend that is nearly flat but very slightly decreasing. This trend is in fact expected, for it indicates that the W7-X coils are just close enough to the plasma to achieve the necessary control of B_n, consistent with the aim of the W7-X designers to push the coils as far from the plasma as possible. If coils were further from the plasma so as to lie on the expanded control surface, the singular values of the inductance and transfer matrices decrease more rapidly (figure 14.c,) yielding feasibility sequences in figures 19.c-d that are generally increasing. The amplitude of B_n or current needed on the control surface is thus unbounded, and one would effectively need to control an infinite number of the distributions on the control surface. Thus, the nearly flat trends in figures 19.c-d for the real control surface are consistent with the choice of this winding surface for the actual W7-X coils: it is as far from the plasma as possible, at limit of where it is practical to achieve total $B_n = 0$ on the desired plasma surface.

VIII. UNIVERSALITY AND UNIQUENESS: CHOICE OF ANGLES AND BASIS FUNCTIONS

We now analyze several technical issues related to the uniqueness of the transfer and inductance matrices: are the SVDs of these matrices independent of the specific basis functions chosen? And are the singular values and vectors independent of the choice of poloidal and toroidal angles θ and ζ? For example, one could parameterize the plasma surface using the coordinates of the `VMEC` code (in which magnetic field lines are not straight), or the straight-field-line coordinates associated with the normal cylindrical toroidal angle, or Boozer coordinates, etc. We will prove and demonstrate numerically that the SVDs are indeed independent of these choices of basis functions and angles, under

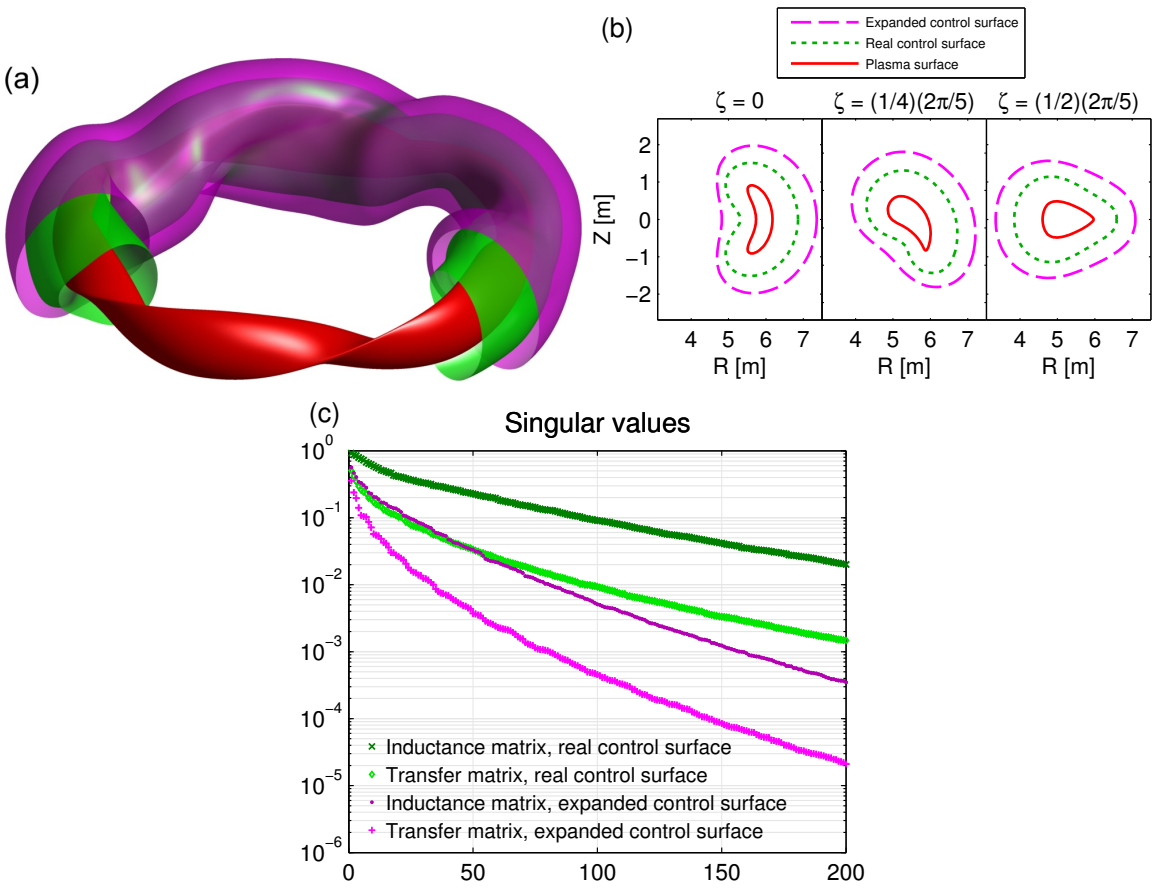


FIG. 14. (Color online) (a) Plasma surface (red) and the two control surfaces considered for the W7-X scenario. (b) Cross-sections of the plasma and control surfaces at three toroidal angles. (c) Singular values of the transfer and inductance matrices. The inductance matrix singular values have been divided by $2\pi^2\mu_0 R_0$ as in previous figures to fit on the same axes. ($R_0 = 5.5$ m.)

reasonable assumptions. Specifically, we loosely prove the following two results for both the inductance and transfer matrices: (1) Two different sets of basis functions which are orthogonal under the same weight w yield the same numerically converged singular values, and the associated singular vectors represent the same patterns in real space; and (2) as long as the weight w is specified in a coordinate-independent manner, changing the choice of angles does not alter the numerically converged singular values or real-space patterns represented by the singular vectors.

A. Change of basis

To carry out the rough proofs, we begin by considering a single surface with two different sets of basis functions, $f_j^{(1)}$ and $f_j^{(2)}$. Let the two basis sets be associated with weights $w^{(1)}$ and $w^{(2)}$ which may or may not be equal. We introduce a matrix $\mathbf{A}$ with components $A_{i,j}$ which relates the bases:

$$f_i^{(2)} = A_{i,j} f_j^{(1)}. \tag{41}$$

(For finite basis sets the basis functions in one set may not be exact linear combinations of the functions in the other set, but once a sufficient number of basis functions are included in each set to resolve a given number of singular vectors to a given level, (41) holds in an approximate sense. Due to this issue, our conclusions will apply only to the singular vectors and values that are well resolved numerically, rather than to all of them.) Applying the operation $\int \mathrm{d}^2a\, w^{(1)}\, f_n^{(1)}(\ldots)$ to (41) gives

$$A_{i,n} = \int \mathrm{d}^2a\, w^{(1)} f_n^{(1)} f_i^{(2)}. \tag{42}$$

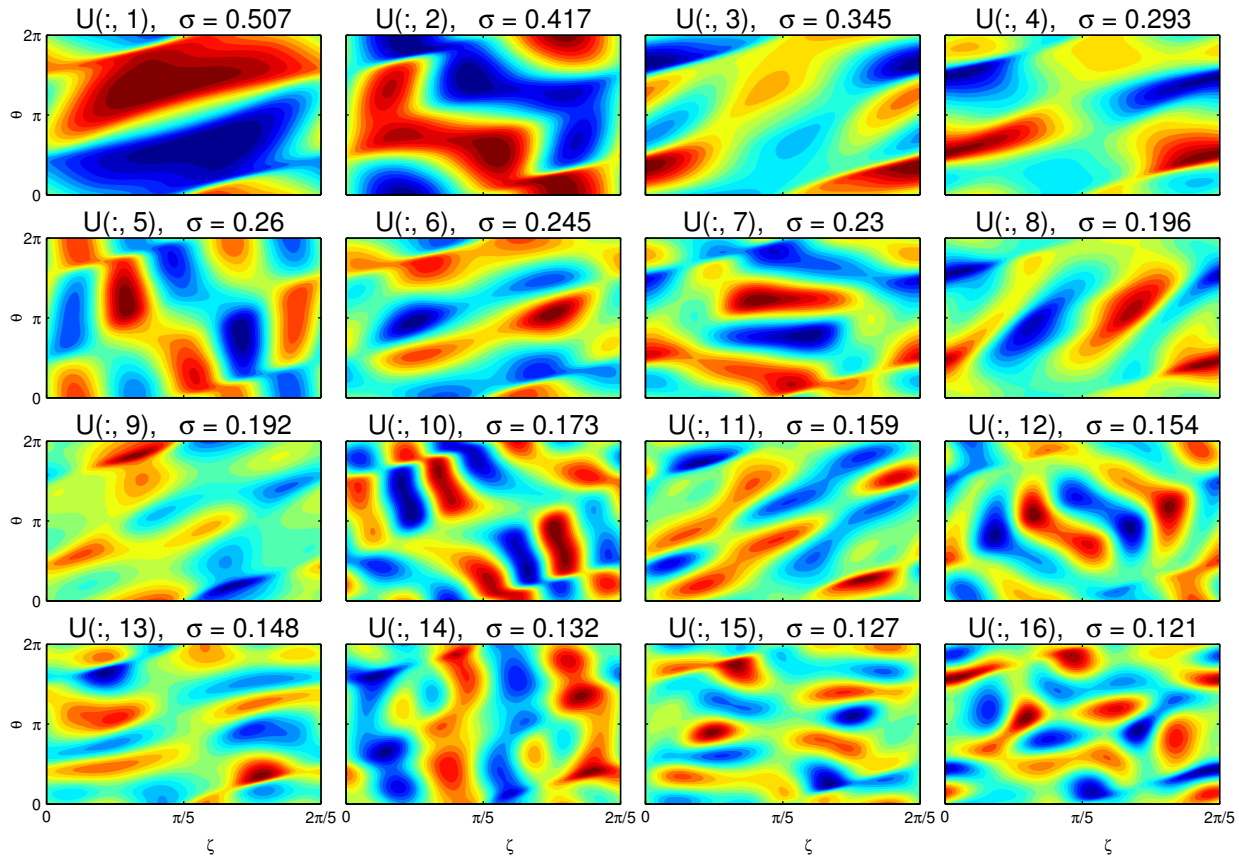


FIG. 15. (Color online) The first 16 left singular vectors $\mathbf{U}$ of the transfer matrix for the W7-X scenario of figure 14, corresponding to distributions of B_n on the plasma surface.

We now form

$$\begin{aligned}\left(\mathbf{A}\mathbf{A}^T\right)_{i,j} = A_{i,k}A_{j,k} &= \int \mathrm{d}^2a\; w^{(1)} f_i^{(2)} f_k^{(1)} A_{j,k} \\ &= \int \mathrm{d}^2a\; w^{(1)} f_i^{(2)} f_j^{(2)},\end{aligned} \tag{43}$$

where we have applied (42) to the first factor of $\mathbf{A}$ and then used (41). If $w^{(1)} = w^{(2)}$, then the right-hand side of (43) becomes $\delta_{i,j}$, so $\mathbf{A}\mathbf{A}^T = \mathbf{I}$. We thus arrive at the important conclusion that if the two bases have the same weight, then $\mathbf{A}$ is an orthogonal matrix.

B. Uniqueness of inductance matrix

Now suppose we have two inductance matrices, both relating the current potential on the control surface to B_n on the plasma surface, but using two different sets of basis functions on the plasma surface, $f_i^{(p1)}$ and $f_i^{(p2)}$:

$$M_{i,j}^{(p1,c)} = \int \mathrm{d}^2r^{(p)} \int \mathrm{d}^2r^{(c)} f_i^{(p1)} f_j^{(c)} F\left(\boldsymbol{r}^{(p)}, \boldsymbol{r}^{(c)}\right) \tag{44}$$

and

$$M_{i,j}^{(p2,c)} = \int \mathrm{d}^2r^{(p)} \int \mathrm{d}^2r^{(c)} f_i^{(p2)} f_j^{(c)} F\left(\boldsymbol{r}^{(p)}, \boldsymbol{r}^{(c)}\right), \tag{45}$$

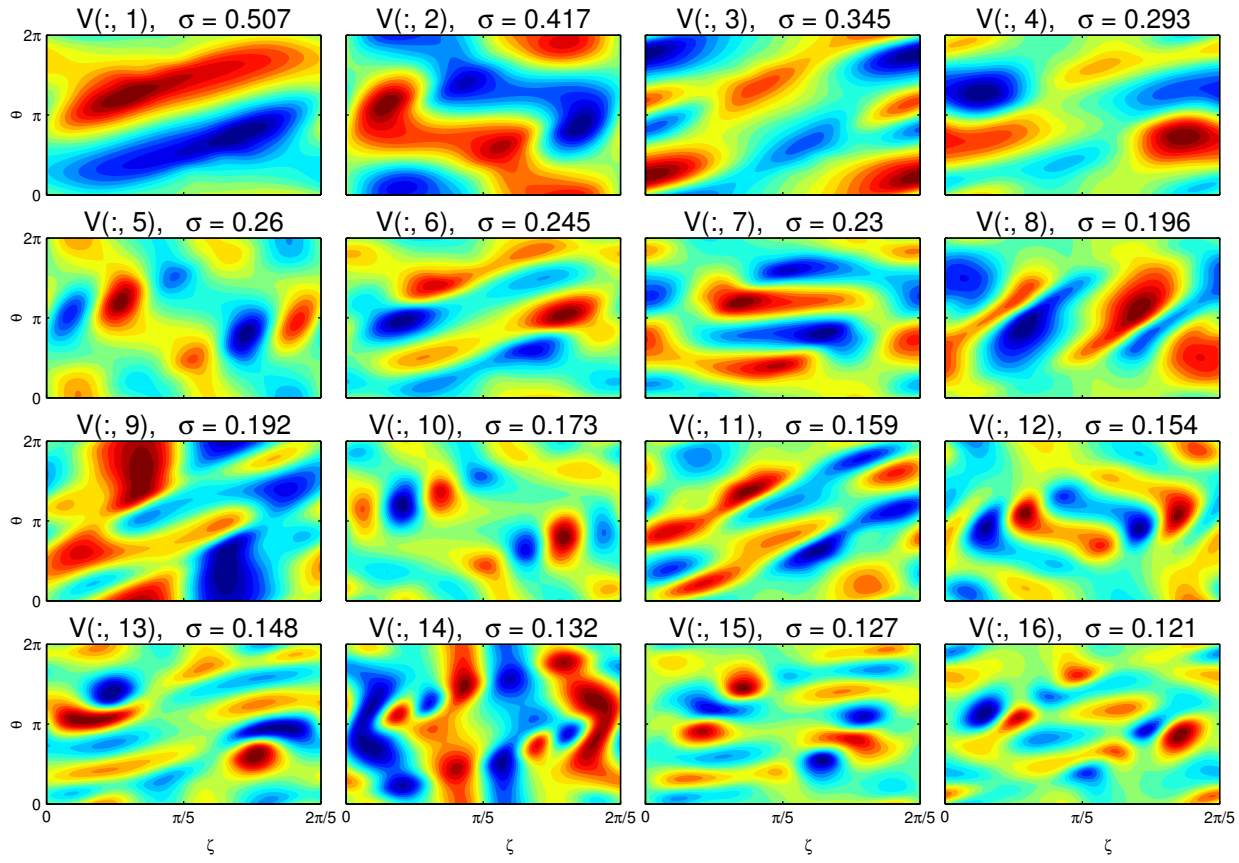


FIG. 16. (Color online) The first 16 right singular vectors $\mathbf{V}$ of the transfer matrix for the W7-X scenario of figure 14, corresponding to distributions of B_n on the control surface.

where the integrals are over the two surfaces, superscripts indicate p = plasma surface and c = control surface, and

$$F\left(\boldsymbol{r}^{(p)},\boldsymbol{r}^{(c)}\right)=\frac{\mu_0}{4\pi}\left[\frac{3\left(\boldsymbol{r}^{(p)}-\boldsymbol{r}^{(c)}\right)\cdot\boldsymbol{n}^{(p)}\left(\boldsymbol{r}^{(p)}-\boldsymbol{r}^{(c)}\right)\cdot\boldsymbol{n}^{(c)}}{\left|\boldsymbol{r}^{(p)}-\boldsymbol{r}^{(c)}\right|^5}-\frac{\boldsymbol{n}^{(p)}\cdot\boldsymbol{n}^{(c)}}{\left|\boldsymbol{r}^{(p)}-\boldsymbol{r}^{(c)}\right|^3}\right]. \tag{46}$$

Using (41), (44), and (45),

$$M_{i,j}^{(p2,c)}=A_{i,k}M_{k,j}^{(p1,c)}. \tag{47}$$

Now if we have the SVD for basis 1,

$$\mathbf{M}^{(p1,c)}=\mathbf{U}^{(1)}\mathbf{\Sigma}^{(1)}\left[\mathbf{V}^{(1)}\right]^T, \tag{48}$$

(47) implies

$$\mathbf{M}^{(p2,c)}=\mathbf{A}\mathbf{U}^{(1)}\mathbf{\Sigma}^{(1)}\left[\mathbf{V}^{(1)}\right]^T. \tag{49}$$

Let us now suppose the two sets of basis functions use the same weight w so the matrix $\mathbf{A}$ is orthogonal, as proved above. Then since both $\mathbf{A}$ and $\mathbf{U}^{(1)}$ are orthogonal, so is $\mathbf{A}\mathbf{U}^{(1)}$. Hence the SVD for basis 2 is

$$\mathbf{M}^{(p2,c)}=\mathbf{U}^{(2)}\mathbf{\Sigma}^{(2)}\left[\mathbf{V}^{(2)}\right]^T \tag{50}$$

where

$$\begin{aligned}\mathbf{U}^{(2)}&=\mathbf{A}\mathbf{U}^{(1)},\\ \mathbf{\Sigma}^{(2)}&=\mathbf{\Sigma}^{(1)},\\ \mathbf{V}^{(2)}&=\mathbf{V}^{(1)}.\end{aligned} \tag{51}$$

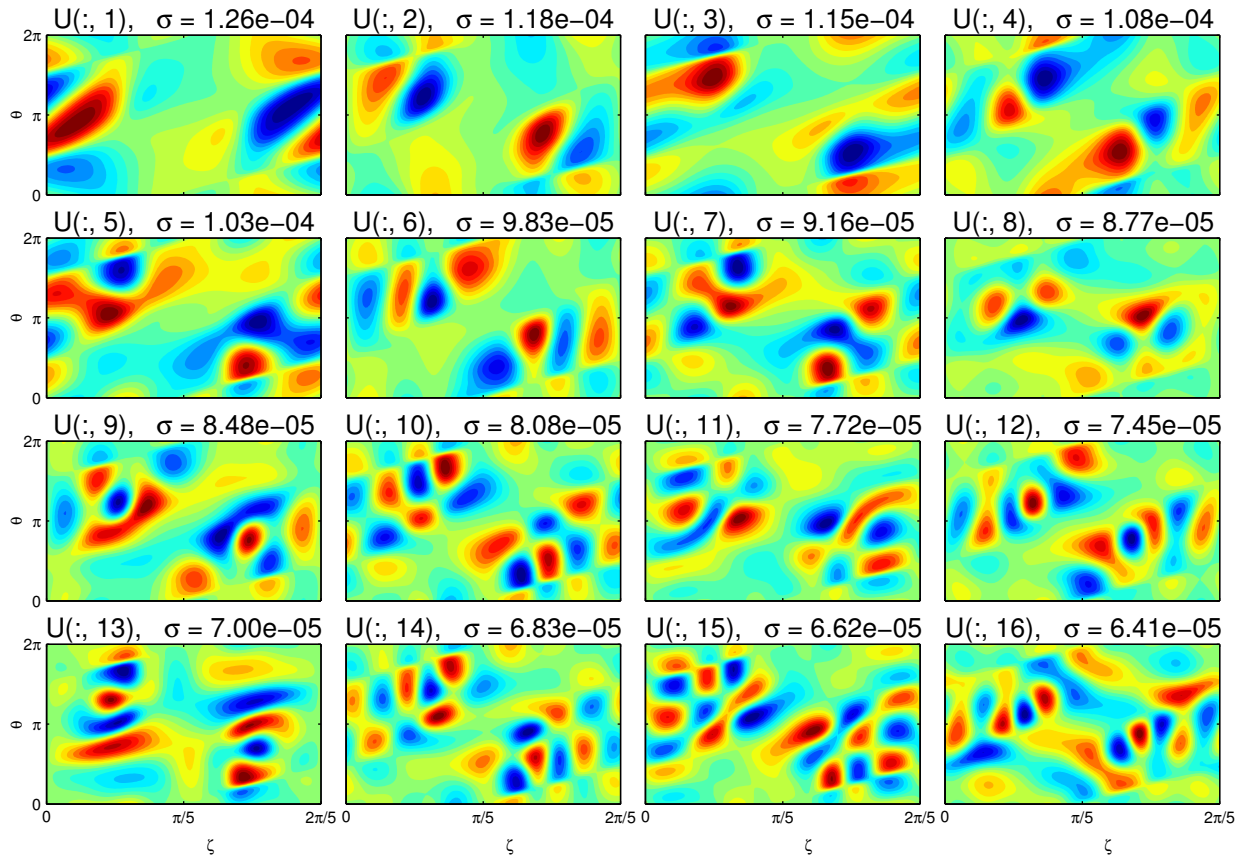


FIG. 17. (Color online) The first 16 left singular vectors $\mathbf{U}$ of the inductance matrix for the W7-X scenario of figure 14, corresponding to distributions of B_n on the plasma surface.

We have thus proved that the singular values and the right singular vectors are unchanged under the change of basis. The ith left singular vector in basis 1 represents the following physical pattern of B_n on the plasma surface:

$$f_j^{(p1)} U_{j,i}^{(1)}. \tag{52}$$

The ith singular vector in basis 2 is

$$f_j^{(p2)} U_{j,i}^{(2)} = f_k^{(p1)} U_{k,i}^{(1)} \tag{53}$$

where we have used (41), then (51), then (43). Thus, although the matrices $\mathbf{U}^{(1)}$ and $\mathbf{U}^{(2)}$ have different matrix elements, the singular vectors represent the same physical pattern of B_n in real space on the plasma surface.

We can repeat the above steps, considering a change of basis on the control surface instead of on the plasma surface. By doing so we obtain the analogous result: the singular values and left singular vectors are unchanged, and the right-singular vectors correspond to the same physical current potential. This completes the proof that the inductance matrix SVD is effectively independent of the choice of basis functions under the given assumptions.

C. Uniqueness of transfer matrix

We now prove the transfer matrix has the same properties proved of the inductance matrix in the previous section. To do so, we recall the discussion in section IV that the transfer matrix can be written in terms of two inductance matrices: introducing a third surface outside the control surface, and considering the outer-to-plasma-surface inductance matrix $\mathbf{M}^{(p,o)}$ and the outer-to-control-surface inductance matrix $\mathbf{M}^{(c,o)}$, the transfer matrix can be computed from $\mathbf{T} = \mathbf{M}^{(p,o)} \left[\mathbf{M}^{(c,o)}\right]^{-1}$. In the previous section regarding the basis-independence of the inductance matrix, the two surfaces used were arbitrary, so the results apply equally to the outer-to-plasma-surface inductance matrix and the outer-to-control-surface inductance matrix.

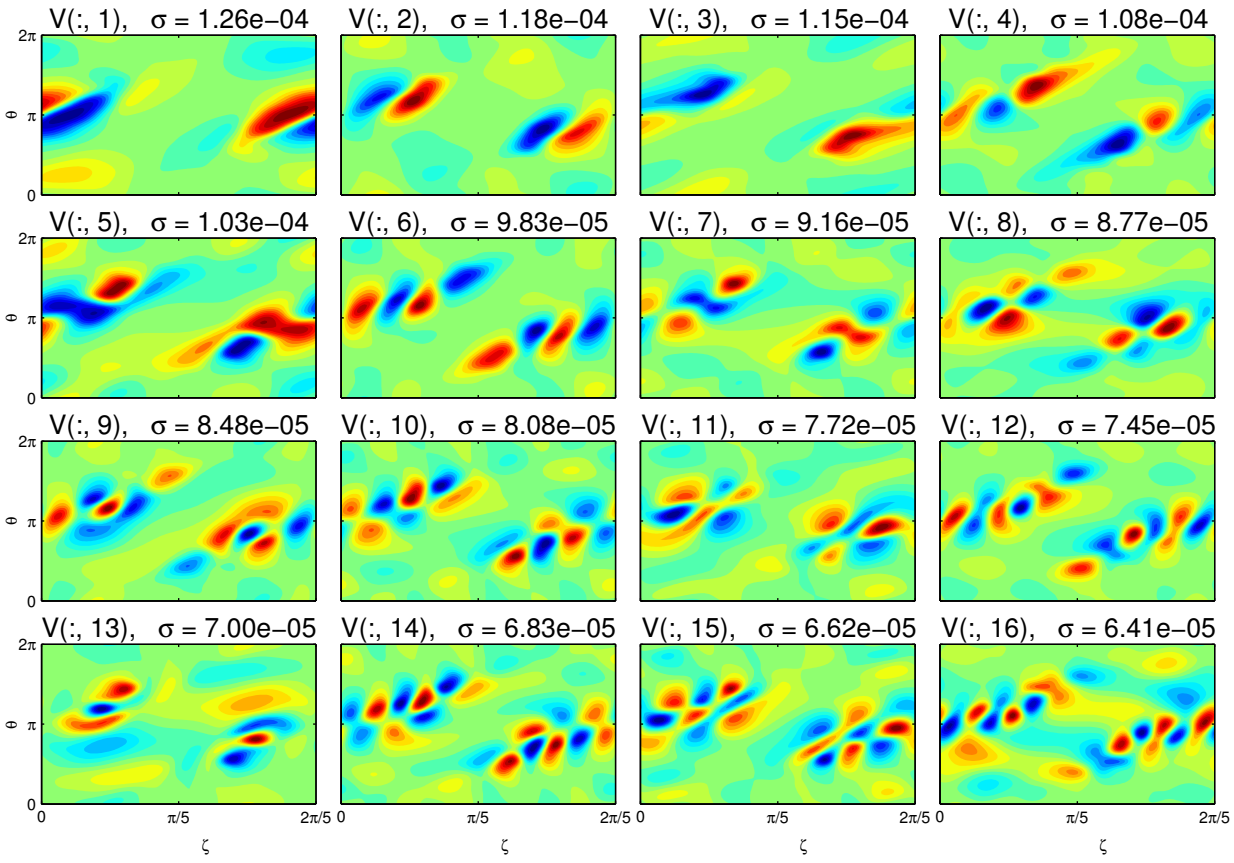


FIG. 18. (Color online) The first 16 right singular vectors $\mathbf{V}$ of the inductance matrix for the W7-X scenario of figure 14, corresponding to distributions of current potential κ on the control surface.

Suppose we consider two basis sets on the plasma surface, orthogonal under the same weight, as in the previous section. Let us denote the two associated inductance matrices between the outer and plasma surfaces by $\mathbf{M}^{(p1,o)}$ and $\mathbf{M}^{(p2,o)}$. From (41), (44), and (45),

$$\mathbf{M}^{(p2,o)} = \mathbf{A}\mathbf{M}^{(p1,o)}. \tag{54}$$

The transfer matrix for basis 1 is

$$\mathbf{T}^{(1)} = \mathbf{M}^{(p1,o)}\left[\mathbf{M}^{(c,o)}\right]^{-1}, \tag{55}$$

and the transfer matrix for basis 2 is

$$\mathbf{T}^{(2)} = \mathbf{M}^{(p2,o)}\left[\mathbf{M}^{(c,o)}\right]^{-1} = \mathbf{A}\mathbf{M}^{(p1,o)}\left[\mathbf{M}^{(c,o)}\right]^{-1} = \mathbf{A}\mathbf{T}^{(1)}. \tag{56}$$

Hence, if the SVDs of the two transfer matrices are

$$\mathbf{T}^{(1)} = \mathbf{U}^{(1)}\mathbf{\Sigma}^{(1)}\left[\mathbf{V}^{(1)}\right]^{T} \tag{57}$$

and

$$\mathbf{T}^{(2)} = \mathbf{U}^{(2)}\mathbf{\Sigma}^{(2)}\left[\mathbf{V}^{(2)}\right]^{T}, \tag{58}$$

the fact that $\mathbf{A}$ is orthogonal in (56) implies

$$\begin{aligned} \mathbf{U}^{(2)} &= \mathbf{A}\mathbf{U}^{(1)}, \\ \mathbf{\Sigma}^{(2)} &= \mathbf{\Sigma}^{(1)}, \\ \mathbf{V}^{(2)} &= \mathbf{V}^{(1)}, \end{aligned} \tag{59}$$

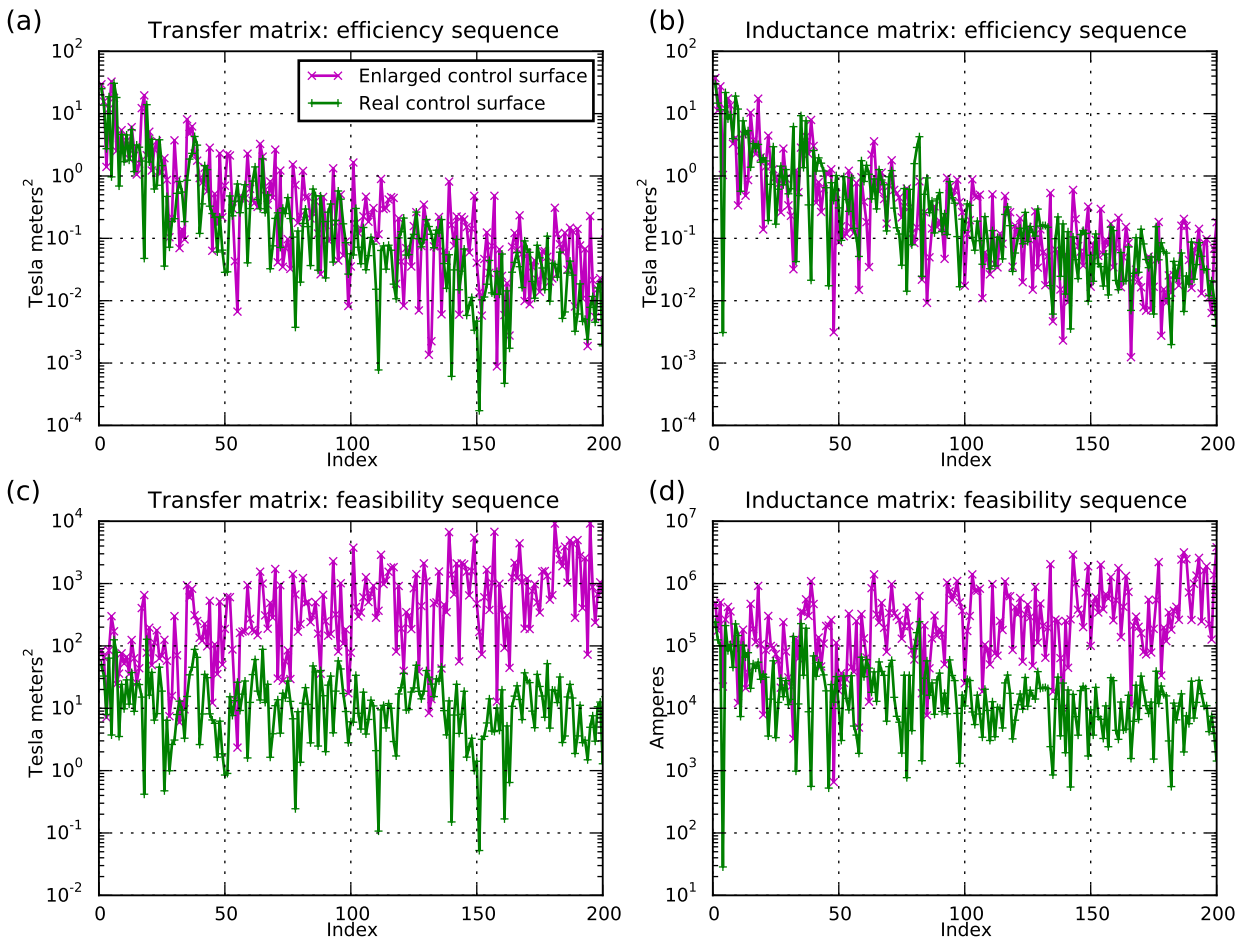


FIG. 19. (Color online) (a)-(b) Efficiency sequences for the transfer and inductance matrices of W7-X, defined as the projection of B_n^{fix} onto the left singular vectors. Here B_n^{fix} is B_n on the plasma surface driven by the $G\zeta/(2\pi)$ term in the current potential (4), the term which generates the main toroidal magnetic field. The sequences in (a)-(b) are divided by the singular values in (c)-(d) to give the required amplitudes of magnetic flux or current on the control surface. The general rate of decrease in the efficiency sequence is insensitive to the control surface, whereas the rate of increase in the feasibility sequence increases with the distance between plasma and control surfaces. The nearly flat trend in (c)-(d) for the real control surface indicates it is difficult to produce the W7-X plasma shape using coils further from the plasma than the actual W7-X coils.

exactly as in the previous section. Thus, by the same reasoning as before, a change of basis on the plasma surface leaves the singular values and right singular vectors unchanged, and the left singular vectors correspond to the same physical pattern of B_n on the plasma surface. Repeating this analysis with a change of basis on the control surface rather than the plasma surface, we obtain an analogous result: a change of basis on the control surface leaves the singular values and left singular vectors unchanged, while the right singular vectors correspond to the same physical B_n on the control surface.

D. Numerical example

To demonstrate the above uniqueness properties of the inductance and transfer matrices, we need an alternative set of basis functions f'_j (different from (15)) which are orthogonal under the same constant weight $w = 1/A$. One way to construct such a basis set is the following procedure. Choose any convenient coordinates θ and ζ to parameterize the surface so Fourier basis functions (14) can be evaluated (though we will use different basis functions in the end). Let $L_{i,j}$ be the yet-unknown linear relation between these Fourier functions f_j and the new basis functions f'_j we seek:

$$f_i = L_{i,j} f'_j. \tag{60}$$

Next, consider the matrix

$$C_{i,j} = \frac{1}{A} \int \mathrm{d}^2 a \; f_i f_j, \tag{61}$$

which can be evaluated numerically for any given surface. Applying (60) twice and then applying (5) for the new basis, we find

$$C_{i,j} = L_{i,k} L_{j,k} \quad \Rightarrow \quad \mathbf{C} = \mathbf{L}\mathbf{L}^T. \tag{62}$$

Thus, we can obtain $\mathbf{L}$ by a (lower) Cholesky decomposition of $\mathbf{C}$. With $\mathbf{L}$ now in hand, the new basis functions are then given by

$$f'_i = (\mathbf{L}^{-1})_{i,j} f_j. \tag{63}$$

These basis functions differ from (15). For example, since $\mathbf{L}^{-1}$ is lower-triangular, the first basis function in the new set f'_1 is a pure Fourier function, so due to the variation of N, this first basis function cannot be a member of the set (15). Note also that the construction above does not give a unique set of new (constant-w) basis functions: changing the order of m_j and n_j in the Fourier functions will give a different $C_{i,j}$ and hence a different set of basis functions f'_j.

Figure 20 numerically demonstrates the independence of the SVDs with respect to the basis set for an NCSX scenario. The geometry is illustrated in figures 20.a-b. The plasma surface is taken from the outer surface of a `VMEC` equilibrium, and we take the control surface to be offset by a uniform distance from the plasma surface. Only stellarator-symmetric distributions are included. Figure 20.c displays the first 50 singular values computed using different choices for the basis functions on the plasma and control surface. The 'standard basis' functions (15) are used unless specified otherwise. Results obtained using the standard basis functions on both surfaces are shown with the $\diamondsuit$ symbol. Using the standard basis functions on the control surface with the 'Cholesky basis' functions (63) on the plasma surface gives the + dataseries. Using the standard basis functions on the plasma surface with the Cholesky basis functions on the control surface gives the $\times$ dataseries. It can be seen that these three methods give identical singular values, as predicted by our proof above. The fact that these quite different basis set constructions yield identical singular values highlights the universality of our SVD analysis. Somewhat different singular values are obtained if the 'pure' Fourier functions (14) are used as the basis functions on either surface, as illustrated with $\bullet$ symbols. As already noted, the pure Fourier functions are only orthogonal under a non-constant weight, i.e. a different weight than the standard or Cholesky basis sets, and so the singular values will generally differ.

A similar comparison of the transfer matrix singular values is shown in Figure 20.d. As with the inductance matrix, and as proved analytically above, any combination of the standard basis and Cholesky basis sets on the plasma and control surface yields identical singular values, since these basis sets share the same orthogonality weight. Use of pure Fourier basis functions results in slightly different singular values.

E. Change of angle coordinates

One important consequence of sections VIII B-VIII C is that the inductance and transfer matrix SVDs are independent (up to discretization error) of the choice of angles used to parameterize the surfaces. Keeping the weight w fixed, consider the basis functions (15) for one pair of (θ, ζ) angles and the basis functions (15) for a different pair of (θ, ζ) angles. To these two sets of basis functions, we can apply the results of sections VIII B-VIII C. Thus, SVD calculations with the two sets of angles must yield the same singular values, and the singular vectors must represent the same patterns in real space.

Note however that the same reasoning does not apply to the Fourier functions (14). These functions are only a valid orthonormal basis if the orthogonality weight is $w = 1/(4\pi^2 N)$, with $N(\theta, \zeta) = |\boldsymbol{N}|$ given by (12), which *changes if we change the* (θ, ζ) *coordinates*. If we alter the coordinates used to parameterize the surfaces, w changes, and so there is no reason the singular values and vectors will not change. Thus, for maximal universality of the results, there is good motivation to use the basis functions (15) (as we have done throughout this work) and not (14).

These points are illustrated in figures 20.e-f. The same plasma surface shape is parameterized by two different poloidal angles: the angle used in the `VMEC` code (chosen to minimize the spectral width of the surface shape), and the PEST angle (which makes magnetic field lines straight). Both coordinate systems use the standard cylindrical toroidal angle, and the coordinates describing the control surface are held fixed. The figures show that when the coordinate-independent weight $w = 1/A$ is used (associated with the standard basis functions), the singular values of the inductance and transfer matrices are independent of the choice of angle. However, when the Fourier basis functions are used (implying $w = 1/(4\pi^2 N)$), the singular values do depend on the choice of angle.

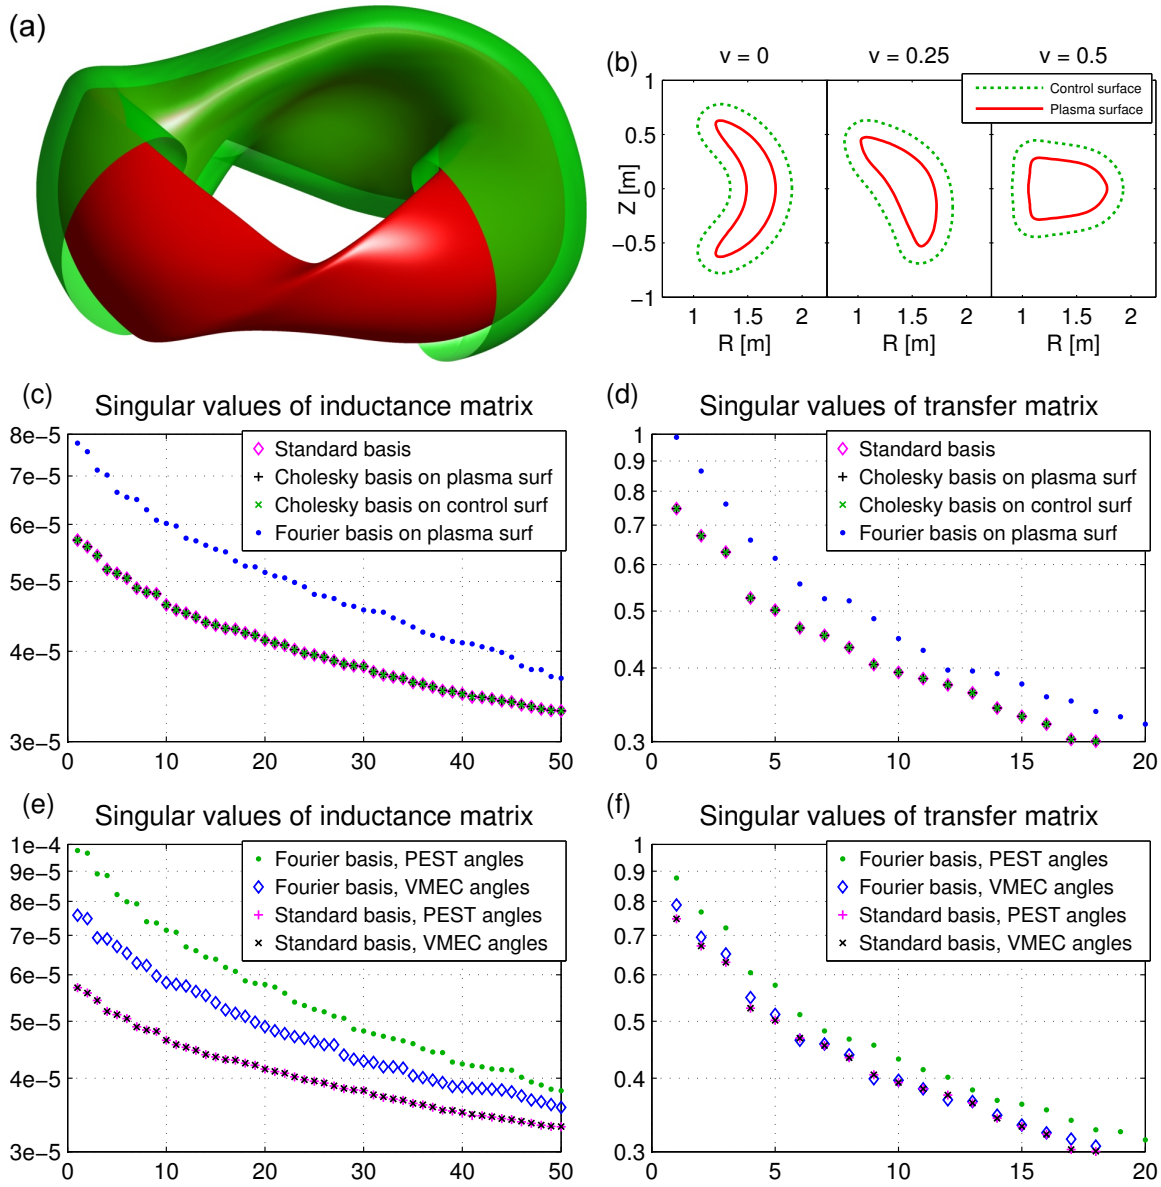


FIG. 20. (Color online) (a) NCSX plasma surface (red) and control surface (green) for demonstrating the uniqueness results. (b) Cross-sections of the plasma and control surfaces at three toroidal angles. (c) The singular values of the inductance matrix are unchanged if different basis functions are used which share the same orthogonality weight ($\diamondsuit, +, \times$), whereas the singular values change somewhat when changing to a basis with different weight ($\bullet$). (d) Same result for the transfer matrix. (e) Singular values of the inductance matrix are independent of the coordinates used to parameterize the surfaces if a coordinate-independent weight is used, as in the 'standard basis' (15), but not if Fourier functions (14) are used. (f) The same is true of the transfer matrix.

IX. DISCUSSION AND CONCLUSIONS

To minimize the engineering difficulties in a tokamak or stellarator reactor, the distance between plasma and coils should be large, and this distance can be much greater for 'efficient' plasma shapes. The transfer and inductance matrices, and the singular value decompositions thereof, provide information as to how efficiently a plasma shape can be produced from a distance. One quantitative measure of difficulty derived from these matrices is the efficiency sequence, the rate of (exponential) decrease in the projection of $B_n^{pl} + B_n^{fix}$ (normal magnetic field due to plasma current and any fixed known external coils) onto the left singular vectors. This measure is rather insensitive to the choice of control surface. Another measure is the feasibility sequence, which is the efficiency sequence divided by the singular values. This sequence decreases or increases if it is likely feasible or infeasible to support the plasma shape

using currents on the control surface.

In this work we have further developed the theory of the transfer and inductance matrices, and demonstrated their properties with several numerical computations. Specifically,

1. In figure 10, we demonstrated that the transfer and inductance matrices both correctly capture the fact that a diverted tokamak shape is not significantly harder to produce than an elliptical shape, despite the sharp corner. The figure also shows that the transfer and inductance matrices capture the fact that concave shapes are inefficient/difficult to produce.

2. For axisymmetric surfaces with circular cross section, we computed the SVDs analytically in the large-aspect-ratio approximation, equations (30)-(31) and (35)-(36). We demonstrated by direct numerical computation that these formulae remain extremely accurate at realistic aspect ratio. These formulae may be applicable for other problems that do not directly involve the transfer or inductance matrices, since they enable pairs of poloidal and toroidal mode numbers (m, n) to be ranked by efficiency.

3. We derived eq (38) relating the toroidal and poloidal mode numbers that can be controlled with comparable efficiency, and we demonstrated this formula agrees well with direct calculation of the transfer matrix singular values in toroidal geometry at realistic aspect ratio.

4. We systematically compared the inductance and transfer matrix approaches. While the two methods generally yield the same trends that shorter spatial wavelengths are less efficient, structures with the longest wavelengths are found to be the most efficient in the transfer matrix approach but not the most efficient in the inductance matrix approach. This behavior arises physically because the inductance matrix relates B_n to the current potential, B_n arises from the *gradient* of the current potential, and this gradient accentuates smaller structures. This difference between the inductance and transfer matrix measures has been visible in all the examples we considered, such as the different $m = 0$ behavior in figure 2, the opposite n dependencies in figure 4.a-b, the more complicated structures in figure 11.d compared to 11.c, and the shorter wavelengths in figure 18 compared to figure 16.

5. We proved analytically and demonstrated numerically that the SVD efficiency measures using both the transfer and inductance matrices are unique, in the following sense. Assuming a coordinate-independent orthogonality weight is used, and assuming a sufficient number of basis functions are included to adequately resolve a given number of singular vectors, then both the resolved singular values and the real-space representation of the resolved singular vectors are independent of the basis functions and independent of the coordinates used to parameterize the surfaces.

We anticipate these developments in the theory of efficient plasma shaping will facilitate stellarator optimization and tokamak control. For example, if future stellarator optimizations target plasma shape efficiency using the above measures, one can be confident that these measures will not erroneously penalize shapes with sharp edges that may be advantageous for a divertor.

ACKNOWLEDGMENTS

We are grateful to Michael Drevlak for providing the W7-X plasma and coil surface data, and to Geri Papp for assistance obtaining ITER data. We also thank Sam Lazerson for assistance with the `VMEC` and `NESCOIL` codes, and David Gates for feedback on the manuscript. This work was supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Science, under Award Numbers DE-FG02-93ER54197, DE-FC02-08ER54964, and DE-FG02-95ER54333. Computations were performed on the Edison and Cori systems at the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Appendix A: Overlap matrix

Here we present one method to visualize the relationship between the inductance and transfer matrices. Given any two measures of efficiency, each associated with an orthonormal efficiency-ordered set of functions on the plasma surface, one illuminating method to compare the measures is to expand one set of functions in the other set. For the inductance and transfer matrix singular vectors, this expansion is given by the "overlap matrix" $\mathbf{U}^T_{\text{transfer}}\mathbf{U}_{\text{inductance}}$.

This overlap matrix is shown in figure 21 for the scenario of figure 1.a, axisymmetric surfaces with circular cross-section and plasma aspect ratio 10. Several facts can be understood from this figure. First, each row and column are dominated by a single $O(1)$ matrix element, which indicates that the two efficiency measures result in essentially the same set of B_n distributions, if one neglects the order. However, these dominant $O(1)$ matrix elements can be rather far from the diagonal, particularly for the first hundred distributions, indicating that the two efficiency measures rank these distributions in a different order. After several hundred distributions, the dominant elements of the overlap matrix become more clustered near the diagonal, indicating the ranking of the two efficiency measures becomes more similar.

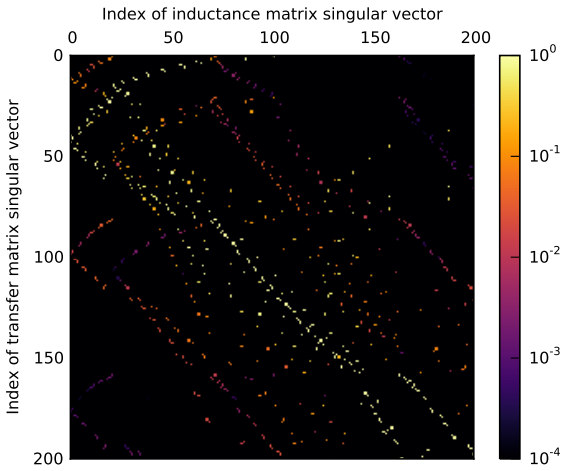


FIG. 21. (Color online) Elements of the overlap matrix $\mathbf{U}^T_{\text{transfer}}\mathbf{U}_{\text{inductance}}$, showing an expansion of the transfer matrix B_n distributions in the inductance matrix B_n distributions, and vice-versa. This calculation is performed for axisymmetric surfaces with circular cross-section, $R/a_P = 10$,and $a_C/a_P = 1.7$ (figure 1.a).

Figure 22 shows the overlap matrix for the divertor shape with the 1.5 m offset control surface, from figure 7. As this overlap matrix is dominated by the diagonal for indices above ~ 10, the inductance and transfer matrices provide essentially identical information beyond the 10th singular vectors. However, the large values of the overlap matrix above the diagonal for indices < 10 result from long-wavelength magnetic field structures that are weighted as more efficient by the transfer matrix than by the inductance matrix. For instance, the 10th inductance matrix singular vector is dominated by an $m = 1$ structure.

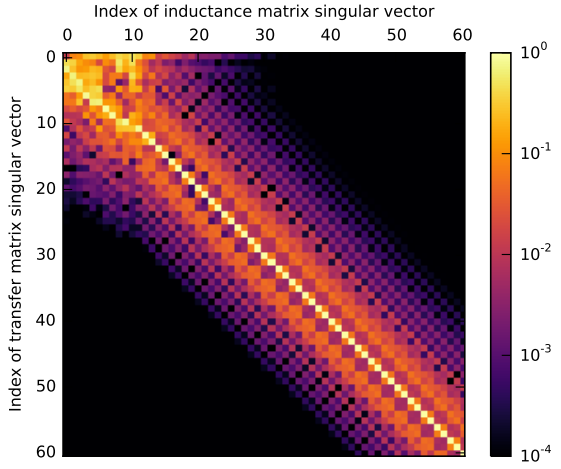


FIG. 22. (Color online) Elements of the overlap matrix $\mathbf{U}^T_{\text{transfer}}\mathbf{U}_{\text{inductance}}$. This calculation is performed for the axisymmetric divertor plasma shape with 1.5 m offset control surface (figure 7). Note that only the axisymmetric singular vectors are included here.

Elements for the overlap matrix $\mathbf{U}^T_{\text{transfer}}\mathbf{U}_{\text{inductance}}$ of the W7-X example are shown in figure 23, again illustrating the transfer matrix singular vectors expanded in the inductance matrix singular vectors, and equivalently the

inductance matrix singular vectors expanded in the transfer matrix singular vectors. It can be seen that the overlap matrix is generally largest near the diagonal, indicating the inductance and transfer matrix give a roughly similar sense of efficiency. However, as one moves away from the diagonal of the overlap matrix, the decrease of the matrix element sizes is slow. A singular vector of the transfer matrix generally has a projection of at least 1% onto as many as ~ 100 singular vectors of the inductance matrix, and vice-versa. Thus, if several hundred singular vectors are being considered, it makes little difference whether one measures efficiency using the transfer or inductance matrix approaches, but if ≤ 100 singular vectors are being considered, it makes a large difference. Figure 23 has one off-diagonal feature, showing a significant overlap between transfer matrix singular vectors of index ~ 9 and inductance matrix singular vectors of index ~ 80. This feature can be understood from the fact that these singular vectors have a significant $(m = 0, n = 1)$ component, which is precisely the component found to be most different between the inductance and transfer measures in section V, efficient according to the transfer matrix but inefficient according to the inductance matrix.

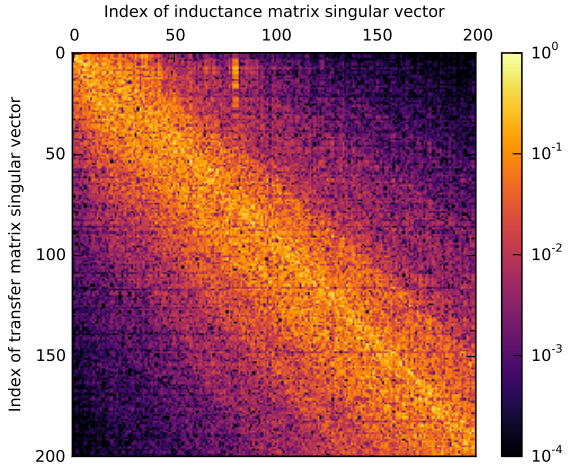


FIG. 23. (Color online) Elements of the overlap matrix $\mathbf{U}^T_{\text{transfer}}\mathbf{U}_{\text{inductance}}$ This calculation is performed for the W7-X geometry of figure 14 with real control surface.

[1] T. Klinger, C. Baylard, C. D. Beidler, J. Boscary, H. S. Bosch, A. Dinklage, D. Hartmann, P. Helander, H. Maassberg, A. Peacock, T. S. Pedersen, T. Rummel, F. Schauer, L. Wegener, and R. Wolf, Fusion Eng. Des. **88**, 461 (2013).
[2] P. Merkel, Nucl. Fusion **27**, 867 (1987).
[3] Note that the condition of small curvature is not simply related to the condition of small spectral width of $R(\theta, \zeta)$ and $Z(\theta, \zeta)$, and vice versa. The circle $R = \cos(\theta)$, $Z = \sin(\theta)$ has low curvature but large spectral width when parameterized in terms of the alternative coordinate $\vartheta = \theta + \sin(\theta)$. Conversely, the shape given by $R = \cos(\theta)$, $Z = \sin(\theta) + \sin(2\theta)/2$ has small spectral width but infinite curvature.
[4] J. Nührenberg and R. Zille, Phys. Lett. A **129**, 113 (1988).
[5] N. Pomphrey, L. Berry, A. Boozer, A. Brooks, R. E. Hatcher, S. P. Hirshman, L.-P. Ku, W. H. Miner, H. E. Mynick, W. Reiersen, D. J. Strickler, and P. M. Valanju, Nucl. Fusion **41**, 339 (2001).
[6] A. Boozer, Phys. Plasmas **7**, 629 (2000).
[7] A. Boozer, Phys. Plasmas **7**, 3378 (2000).
[8] A. Boozer, Nucl. Fusion **55**, 025001 (2015).
[9] A. Boozer, J. Plasma Phys. **81**, 515810606 (2015).
[10] S. P. Hirshman and J. C. Whitson, Phys. Fluids **26**, 3553 (1983).
[11] J. L. Hess, Comp. Meth. Appl. Mech. Engineering **5**, 145 (1975).
[12] M. A. Jaswon and G. T. Symm, *Integral equation methods in potential theory and electrostatics* (Academic Press, 1977).
[13] P. Merkel, J. Comp. Phys. **66**, 83 (1986).
[14] P. Merkel, in *Proceedings of the 1987 Workshop on the Theory of Fusion Plasmas*, Vol. 25 (Varenna, Italy, 1988) p. 25.
[15] A. A. Becker, *The bounday element method in engineering* (McGraw-Hill, 1992).
[16] W. Dommaschk, Comp. Phys. Comm. **40**, 203 (1986).
[17] V. D. Shafranov and L. E. Zakharov, Nucl. Fusion **12**, 599 (1972).
[18] M. Drevlak, D. Monticello, and A. Reiman, Nucl. Fusion **45**, 731 (2005).